Noname manuscript No.
(will be inserted by the editor)

Cathoristic logic

A modal logic of incompatible propositions

Richard Prideaux Evans · Martin Berger

arXiv:1411.7158v1 [cs.LO] 26 Nov 2014

Received: date / Accepted: date

Keywords Modal logic, Hennessy-Milner logic, transition systems, negation, exclusion, elementary equivalence, incompatibility semantics, knowledge representation, philosophy of language.

Richard Prideaux Evans, Imperial College E-mail: richardprideauxevans@imperial.ac.uk · Martin Berger, University of Sussex. E-mail: M.F.Berger@sussex.ac.uk.

Abstract Cathoristic logic is a multi-modal logic where negation is replaced by a novel operator allowing the expression of incompatible sentences. We present the syntax and semantics of the logic including complete proof rules, and establish a number of results such as compactness, a semantic characterisation of elementary equivalence, the existence of a quadratic-time decision procedure, and Brandom's incompatibility semantics property. We demonstrate the usefulness of the logic as a language for knowledge representation.

Contents

1 Introduction . . . 5
1.1 Material incompatibility and negation . . . 6
1.2 Negation as the minimal incompatible . . . 8
1.3 Inferences between atomic sentences . . . 9
1.4 Wittgenstein's vision of a logic of elementary propositions . . . 10
1.5 Outline . . . 10
2 Mathematical preliminaries . . . 12
3 Cathoristic logic . . . 14
3.1 Syntax . . . 14
3.2 Semantics . . . 14
4 Inferences between atomic sentences . . . 18
4.1 Intra-atomic inferences in cathoristic logic . . . 18
4.2 Intra-atomic inferences in first-order logic . . . 20
5 Cathoristic logic as a language for knowledge representation . . . 22
5.1 Representing facts in cathoristic logic . . . 22
5.2 Simpler postconditions . . . 24
5.3 Using tantum ! to optimise preconditions . . . 24
6 Semantics and Decision Procedure . . . 26
6.1 Semantic characterisation of elementary equivalence . . . 26
6.2 Quotienting models . . . 28
6.3 The bounded lattice of models . . . 29
6.4 Computing the least upper bound of the models that satisfy a formula . . . 30
6.5 A decision procedure for cathoristic logic . . . 36
6.6 Incompatibility semantics . . . 37
7 Inference Rules . . . 40
7.1 Example inferences . . . 41
7.2 !-Left and !-Right . . . 41
7.3 Characteristic formulae . . . 41
7.4 Soundness and completeness . . . 43
7.5 Proofs of Lemmas 8, 9 and 10 . . . 44
8 Compactness and the standard translation to first-order logic . . . 49
8.1 Translating from cathoristic to first-order logic . . . 49
8.2 Compactness by translation . . . 51
9 Cathoristic logic and negation . . . 56
9.1 Syntax and semantics . . . 56
9.2 Decision procedure . . . 57
10 Quantified cathoristic logic . . . 59
11 Related work . . . 61
11.1 Brandom's incompatibility semantics . . . 61
11.2 Peregrin on defining a negation operator . . . 62
11.3 Peregrin and Turbanti on defining a necessity operator . . . 63
11.4 Linear logic . . . 63
11.5 Process calculus . . . 64
11.6 Linguistics . . . 64
12 Open problems . . . 65
12.1 Excluded middle . . . 65
12.2 Understanding the expressive strength of cathoristic logic . . . 65
12.3 Acknowledgements . . . 69
A Alternative semantics for cathoristic logic . . . 71
A.1 Pure cathoristic models . . . 71
A.2 Relationship between pure and cathoristic models . . . 71
A.3 Non-determinism and cathoristic models . . . 72
A.4 Semantic characterisation of elementary equivalence . . . 73
B Omitted proofs . . . 76
B.1 Proof of Lemma 5 . . . 76

B.2 Proof of Lemma 6 76
B.3 Proof of Lemma 7 77

1 Introduction

Natural language is full of incompatible alternatives. If Pierre is the current king of France, then nobody else can simultaneously fill that role. A traffic light can be green, amber or red - but it cannot be more than one colour at a time. Mutual exclusion is a natural and ubiquitous concept.

First-order logic can represent mutually exclusive alternatives, of course. To say that Pierre is the only king of France, we can write, following Russell:

$$king(france, pierre) \wedge \forall x.(king(france, x) \rightarrow x = pierre).$$

To say that a particular traffic light, tl, is red - and red is its only colour - we could write:

$$colour(tl, red) \wedge \forall x.colour(tl, x) \rightarrow x = red.$$

In this approach, incompatibility is a *derived* concept, reduced to a combination of universal quantification and identity. First-order logic, in other words, uses relatively complex machinery to express a simple concept:

- Quantification's complexity comes from the rules governing the distinction between free and bound variables[1].
- Identity's complexity comes from the infinite collection of axioms required to formalise the indiscernibility of identicals.

The costs of quantification and identity, such as a larger proof search space, have to be borne every time one expresses a sentence that excludes others - even though incompatibility does not, prima facie, appear to have anything to do with the free/bound variable distinction, or require the full power of the identity relation.

This paper introduces an alternative approach, where exclusion is expressed directly, as a first-class concept. Cathoristic logic[2] is the simplest logic we could find in which incompatible statements can be expressed. It is a multi-modal logic, a variant of Hennessy-Milner logic, that replaces negation with a new logical primitive

$$!A$$

pronounced *tantum*[3] A. Here A is a finite set of alternatives, and $!A$ says that the alternatives in A exhaust all possibilities. For example:

$$!\{green, amber, red\}$$

states that nothing but *green*, *amber* or *red* is possible. Our logic uses modalities to state facts, for example $\langle amber \rangle$ expresses that *amber* is currently the

[1] Efficient handling of free/bound variables is an active field of research, e.g. nominal approaches to logic [23]. The problem was put in focus in recent years with the rising interest in the computational cost of syntax manipulation in languages with binders.

[2] "Cathoristic" comes from the Greek $\kappa\alpha\theta o\rho\acute{\iota}\zeta\epsilon\iota\nu$: to impose narrow boundaries. We are grateful to Tim Whitmarsh for suggesting this word.

[3] "Tantum" is Latin for "only".

case. The power of the logic comes from the conjunction of modalities and tantum. For example

$$\langle amber \rangle \wedge !\{green, amber, red\}$$

expresses that *amber* is currently the case and *red* as well as *green* are the only two possible alternatives to *amber*. Any statement that exceeds what tantum A allows, like

$$\langle blue \rangle \wedge !\{green, amber, red\},$$

is necessarily false. When the only options are green, amber, or red, then blue is not permissible. Now to say that Pierre is the only king of France, we write:

$$\langle king \rangle \langle france \rangle (\langle pierre \rangle \wedge !\{pierre\}).$$

Crucially, cathoristic logic's representation involves no universal quantifier and no identity relation. It is a purely propositional formulation. To say that the traffic light is currently red, and red is its only colour, we write:

$$\langle tl \rangle \langle colour \rangle (\langle red \rangle \wedge !\{red\}).$$

This is simpler, both in terms of representation length and computational complexity, than the formulation in first-order logic given on the previous page. Properties changing over time can be expressed by adding extra modalities that can be understood as time-stamps. To say that that the traffic light was red at time t_1 and amber at time t_2, we can write:

$$\langle tl \rangle \langle colour \rangle (\langle t_1 \rangle (\langle red \rangle \wedge !\{red\}) \wedge \langle t_2 \rangle (\langle amber \rangle \wedge !\{amber\}))$$

Change over time can be expressed in first-order logic with bounded quantification - but modalities are succinct and avoid introducing bound variables.

Having claimed that incompatibility is a natural logical concept, not easily expressed in first-order logic[4], we will now argue the following:

- Incompatibility is conceptually prior to negation.
- Negation arises as the weakest form of incompatibility.

1.1 Material incompatibility and negation

Every English speaker knows that

"Jack is male" is incompatible with "Jack is female"

But *why* are these sentences incompatible? The orthodox position is that these sentences are incompatible because of the following general law:

[4] We will precisify this claim in later sections; (1) first-order logic's representation of incompatibility is longer in terms of formula length than cathoristic logic's (see Section 4.2.1); and (2) logic programs in cathoristic logic can be optimised to run significantly faster than their equivalent in first-order logic (see Section 5.3).

If someone is male, then it is not the case that they are female

Recast in first-order logic:

$$\forall x.(male(x) \rightarrow \neg female(x)).$$

In other words, according to the orthodox position, the incompatibility between the two particular sentences depends on a general law involving universal quantification, implication and negation.

Brandom [7] follows Sellars in proposing an alternative explanation: "Jack is male" is incompatible with "Jack is female" because "is male" and "is female" are *materially incompatible* predicates. They claim we can understand incompatible predicates even if we do not understand universal quantification or negation. Material incompatibility is conceptually prior to logical negation.

Imagine, to make this vivid, a primitive people speaking a primordial language of atomic sentences[5]. These people can express sentences that *are* incompatible. But they cannot express *that* they are incompatible. They recognise when atomic sentences are incompatible, and see that one sentence entails another - but their behaviour outreaches their ability to articulate it.

Over time, these people *may* advance to a more sophisticated language where incompatibilities are made explicit, using a negation operator - but this is a later (and optional) development:

> [If negation is added to the language], it lets one say that two claims are materially incompatible:"If a monochromatic patch is red, then it is not blue." That is, negation lets one make explicit in the form of claims - something that can be said and (so) thought - a relation that otherwise remained implicit in what one practically did, namely treat two claims as materially incompatible[6].

But before making this optional explicating step, our primitive people understand incompatibility without understanding negation. If this picture of our primordial language is coherent, then material incompatibility is conceptually independent of logical negation.

Now imagine a modification of our primitive linguistic practice in which no sentences are ever treated as incompatible. If one person says "Jack is male" and another says "Jack is female", nobody counts these claims as *conflicting*. The native speakers never disagree, back down, retract their claims, or justify them. They just say things. Without an understanding of incompatibility, and the variety of behaviour that it engenders, we submit (following Brandom) that there is insufficient richness in the linguistic practice for their sounds to count as assertions. Without material incompatibility, their sounds are just *barks*.

[5] In this paper, we define a sentence as *atomic* if it does not contain another sentence as a syntactic constituent.

[6] [8] pp.47-48

> Suppose the reporter's differential responsive dispositions to call things red are matched by those of a parrot trained to utter the same noises under the same stimulation. What practical capacities of the human distinguish the reporter from the parrot? What, besides the exercise of regular differential responsive dispositions, must one be able to *do*, in order to count as having or grasping *concepts*? ... To grasp or understand a concept is, according to Sellars, to have practical mastery over the inferences it is involved in... The parrot does not treat "That's red" as incompatible with "That's green"[7].

If this claim is also accepted, then material incompatibility is not just conceptually *independent* of logical negation, but conceptually *prior*.

1.2 Negation as the minimal incompatible

In [7] and [8], Brandom describes logical negation as a limiting form of material incompatibility:

> Incompatible sentences are Aristotelian *contraries*. A sentence and its negation are *contradictories*. What is the relation between these? Well, the contradictory is a contrary: any sentence is incompatible with its negation. What distinguishes the contradictory of a sentence from all the rest of its contraries? The contradictory is the *minimal* contrary: the one that is entailed by all the rest. Thus every contrary of "Plane figure f is a circle" - for instance "f is a triangle", "f is an octagon", and so on - entails "f is *not* a circle".

If someone asserts that it is not the case that Pierre is the (only) King of France, we have said very little. There are so many different ways in which it could be true:

- The King of France might be Jacques
- The King of France might be Louis
- ...
- There may be no King of France at all
- There may be no country denoted by the word "France"

Each of these concrete propositions is incompatible with Pierre being the King of France. To say "It is not the case that the King of France is Pierre" is just to claim that one of these indefinitely many concrete possibilities is true. Negation is just the logically weakest form of incompatibility.

In the rest of this paper, we assume - without further argument - that material incompatibility is conceptually prior to logical negation. We develop a simple modal logic to articulate Brandom's intuition: a language, without negation, in which we can nevertheless make incompatible claims.

[7] [7] pp.88-89, our emphasis.

1.3 Inferences between atomic sentences

So far, we have justified the claim that incompatibility is a fundamental logical concept by arguing that incompatibility is conceptually prior to negation. Now incompatibility is an inferential relation between *atomic sentences*. In this subsection, we shall describe *other* inferential relations between atomic sentences - inferential relations that first-order logic cannot articulate (or can only do so awkwardly), but that cathoristic logic handles naturally.

The *atomic sentences* of a natural language can be characterised as the sentences which do not contain any other sentences as constituent parts[8]. According to this criterion, the following are atomic:

- Jack is male
- Jack loves Jill

The following is not atomic:

Jack is male and Jill is female

because it contains the complete sentence "Jack is male" as a syntactic constituent. Note that, according to this criterion, the following *is* atomic, despite using "and":

Jack loves Jill and Joan

Here, "Jack loves Jill" is not a syntactic constituent[9].

There are many types of inferential relations between atomic sentences of a natural language. For example:

- "Jack is male" is incompatible with "Jack is female"
- "Jack loves Jill" implies "Jack loves"
- "Jack walks slowly" implies "Jack walks"
- "Jack loves Jill and Joan" implies "Jack loves Jill"
- "Jack is wet and cold" implies "Jack is cold"

The first of these examples involves an incompatibility relation, while the others involve entailment relations. A key question this paper seeks to answer is: what is the simplest logic that can capture these inferential relations between atomic sentences?

[8] Compare Russell [24] p.117: "A sentence is of atomic form when it contains no logical words and no subordinate sentence". We use a broader notion of atomicity by focusing solely on whether or not it contains a subordinate sentence, allowing logical words such as "and" *as long as they are conjoining noun-phrases* and not sentences.

[9] To see that "Jack loves Jill" is not a constituent of "Jack loves Jill and Joan", observe that "and" conjoins constituents of the *same syntactic type*. But "Jack loves Jill" is a sentence, while "Joan" is a noun. Hence the correct parsing is "Jack (loves (Jill and Joan))", rather than "(Jack loves Jill) and Joan".

1.4 Wittgenstein's vision of a logic of elementary propositions

In the *Tractatus* [34], Wittgenstein claims that the world is a set of atomic sentences in an idealised logical language. Each atomic sentence was supposed to be *logically independent* of every other, so that they could be combined together in every possible permutation, without worrying about their mutual compatibility. But already there were doubts and problem cases. He was aware that certain statements seemed atomic, but did not seem logically independent:

> For two colours, e.g., to be at one place in the visual field is impossible, and indeed logically impossible, for it is excluded by the logical structure of colour. (6.3751)

At the time of writing the *Tractatus*, he hoped that further analysis would reveal that these statements were not really atomic.

Later, in the *Philosophical Remarks* [33], he renounced the thesis of the logical independence of atomic propositions. In §76, talking about incompatible colour predicates, he writes:

> That makes it look as if *a construction might be possible within the elementary proposition.* That is to say, as if there were a construction in logic which didn't work by means of truth functions. What's more, it also seems that these constructions have an effect on one proposition's following logically from another. For, if different degrees exclude one another it follows from the presence of one that the other is not present. In that case, *two elementary propositions can contradict one another.*

Here, he is clearly imagining a logical language in which there are incompatibilities between atomic propositions. In §82:

> This is how it is, what I said in the Tractatus doesn't exhaust the grammatical rules for 'and', 'not', 'or', etc; *there are rules for the truth functions which also deal with the elementary part of the proposition.* The fact that one measurement is right *automatically* excludes all others.

Wittgenstein does not, unfortunately, show us what this language would look like. In this paper, we present cathoristic logic as one way of formalising inferences between atomic sentences.

1.5 Outline

The rest of this paper is organised as follows: The next section briefly recapitulates the mathematical background of our work. Section 3 introduces the syntax and semantics of cathoristic logic with examples. Section 4 discusses how cathoristic logic can be used to model inferences between atomic sentences. Section 5 describes informally how our logic is useful as a knowledge

representation language. Section 6 presents core results of the paper, in particular a semantic characterisation of elementary equivalence and a decision procedure with quadratic time-complexity. The decision procedure has been implemented in Haskell and is available for public use [1] under a liberal open-source license. This section also shows that Brandom's incompatibility semantics condition holds for cathoristic logic. Section 7 presents the proof rules for cathoristic logic and proves completeness. Section 8 provides two translations from cathoristic logic into first-order logic, and proves compactness using one of them. Section 9 investigates a variant of cathoristic logic with an additional negation operator, and provides a decision procedure for this extension that has an exponential time-complexity. Section 10 extends cathoristic logic with first-order quantifiers and sketches the translation of first-order formulae into first-order cathoristic logic. The conclusion surveys related work and lists open problems. Appendix A outlines a different approach to giving the semantics of cathoristic logic, including a characterisation of the corresponding elementary equivalence. The appendix also discusses the question of non-deterministic models. The remaining appendices present routine proof of facts used in the main section.

A reader not interested in mathematical detail is recommended to look only at Chapters 1, 3, 4, 5, the beginning of Chapter 7, and the Conclusion.

2 Mathematical preliminaries

This section briefly surveys the mathematical background of our paper. A fuller account of order-theory can be found in [9]. Labelled transition systems are explored in [16,26] and bisimulations in [25]. Finally, [11] is one of many books on first-order logic.

Order-theory. A *preorder* is a pair $(S, \sqsubseteq)$ where S is a set, and $\sqsubseteq$ is a binary relation on S that is reflexive and transitive. Let $T \subseteq S$ and $x \in S$. We say x is an *upper bound* of T provided $t \sqsubseteq x$ for all $t \in T$. If in addition $x \sqsubseteq y$ for all upper bounds y of T, we say that x is the *least* upper bound of T. The set of all least upper bounds of T is denoted $\bigsqcup T$. *Lower bounds*, *greatest lower bounds* and $\bigsqcap T$ are defined mutatis mutandis. A *partial order* is a preorder $\sqsubseteq$ that is also anti-symmetric. A partial order $(S, \sqsubseteq)$ is a *lattice* if every pair of elements in S has a least upper and a greatest lower bound. A lattice is a *bounded lattice* if it has top and bottom elements $\top$ and $\bot$ such that for all $x \in S$:

$$x \sqcap \bot = \bot \qquad x \sqcup \bot = x \qquad x \sqcap \top = x \qquad x \sqcup \top = \top.$$

If $(S, \sqsubseteq)$ is a preorder, we can turn it into a partial-order by quotienting: let $a \simeq b$ iff $a \sqsubseteq b$ as well as $b \sqsubseteq a$. Clearly $\simeq$ is an equivalence. Let E be the set of all $\simeq$-equivalence classes of S. We get a canonical partial order, denoted $\sqsubseteq_E$, on E by setting: $[a]_\simeq \sqsubseteq_E [b]_\simeq$ whenever $a \sqsubseteq b$. If all relevant upper and lower bounds exist in $(S, \sqsubseteq)$, then $(E, \sqsubseteq_E)$ becomes a bounded lattice by setting

$$[x]_\simeq \sqcap [y]_\simeq = [x \sqcap y]_\simeq \quad [x]_\simeq \sqcup [y]_\simeq = [x \sqcup y]_\simeq \quad \bot_E = [\bot]_\simeq \quad \top_E = [\top]_\simeq.$$

Transition systems. Let Σ be a set of *actions*. A *labelled transition system over* Σ is a pair $(\mathcal{S}, \rightarrow)$ where $\mathcal{S}$ is a set of *states* and $\rightarrow \subseteq \mathcal{S} \times \Sigma \times \mathcal{S}$ is the *transition relation*. We write $x \xrightarrow{a} y$ to abbreviate $(x, a, y) \in \rightarrow$. We let $s, t, w, w', x, y, z, ...$ range over states, $a, a', b, ...$ range over actions and $\mathcal{L}, \mathcal{L}', ...$ range over labelled transition systems. We usually speak of labelled transition systems when the set of actions is clear from the context. We say $\mathcal{L}$ is *deterministic* if $x \xrightarrow{a} y$ and $x \xrightarrow{a} z$ imply that $y = z$. Otherwise $\mathcal{L}$ is *non-deterministic*. A labelled transition system is *finitely branching* if for each state s, the set $\{t \mid s \xrightarrow{a} t\}$ is finite.

Simulations and bisimulations. Given two labelled transition systems $\mathcal{L}_i = (S_i, \rightarrow_i)$ over Σ for $i = 1, 2$, a *simulation from* $\mathcal{L}_1$ *to* $\mathcal{L}_2$ is a relation $\mathcal{R} \subseteq S_1 \times S_2$ such that whenever $(s, s') \in \mathcal{R}$: if $s \xrightarrow{a}_1 s'$ then there exists a transition $t \xrightarrow{a}_2 t'$ with $(t, t') \in \mathcal{R}$. We write $s \preceq_{sim} t$ whenever $(s, t) \in \mathcal{R}$ for some simulation $\mathcal{R}$. We say $\mathcal{R}$ is a *bisimulation between* $\mathcal{L}_1$ *and* $\mathcal{L}_2$ if both, $\mathcal{R}$ and $\mathcal{R}^{-1}$ are simulations. Here $\mathcal{R}^{-1} = \{(y, x) \mid (x, y) \in \mathcal{R}\}$. We say two states s, s' are *bisimilar*, written $s \sim s'$ if there is a bisimulation $\mathcal{R}$ with $(s, s') \in \mathcal{R}$.

First-order logic. A *many-sorted first-order signature* is specified by the following data. A non-empty set of *sorts*, a set *function symbols* with associated *arities*, i.e. non-empty list of sorts $\#(f)$ for each function symbol f; a set

of *relation symbols* with associated *arities*, i.e. a list of sorts $\#(R)$ for each relation symbol R; a set of *constant symbols* with associated *arity*, i.e. a sort $\#(c)$ for each constant symbol c. We say a function symbol f is *n-ary* if $\#(f)$ has length $n+1$. Likewise, a relation symbol is *n-ary* if $\#(R)$ has length n.

Let $\mathcal{S}$ be a signature. An *$\mathcal{S}$-model* $\mathcal{M}$ is an object with the following components. For each sort σ a set U_σ called *universe* of sort σ. The members of U_σ are called *σ-elements* of $\mathcal{M}$; an element $c^{\mathcal{M}}$ of U_σ for each constant c of sort σ; a function $f^{\mathcal{M}} : (U_{\sigma_1} \times \cdots \times U_{\sigma_n}) \to U_\sigma$ for each function symbol f of arity $(\sigma_1, ..., \sigma_n, \sigma)$; a relation $R^{\mathcal{M}} \subseteq U_{\sigma_1} \times \cdots \times U_{\sigma_n}$ for each relation symbol R of arity $(\sigma_1, ..., \sigma_n)$.

Given an infinite set of variables for each sort σ, the *terms* and *first-order formulae* for $\mathcal{S}$ are given by the following grammar

$$\begin{aligned} t &::= x \mid c \mid f(t_1, ..., t_n) \\ \phi &::= t = t' \mid R(t_1, ..., t_n) \mid \neg\phi \mid \phi \wedge \psi \mid \forall x.A \end{aligned}$$

Here x ranges over variables of all sorts, c over constants, R over n-ary relational symbols and f over n-ary function symbols from $\mathcal{S}$. Other logical constructs such as disjunction or existential quantification are given by de Morgan duality, and truth $\top$ is an abbreviation for $x = x$. If $\mathcal{S}$ has just a single sort, we speak of *single-sorted first-order logic* or just *first-order logic*.

Given an $\mathcal{S}$-model $\mathcal{M}$, an *environment*, ranged over by σ, is a partial function from variables to $\mathcal{M}$'s universes. We write $x \mapsto u$ for the environment that maps x to u and is undefined for all other variables. Moreover, if $\sigma, x \mapsto u$ is the environment that is exactly like σ, except that it also maps x to u, assuming that x is not in the domain of σ. The *interpretation* $[\![t]\!]_{\mathcal{M},\sigma}$ of a term t w.r.t. $\mathcal{M}$ and σ is given by the following clauses, assuming that the domain of σ contains all free variables of t:

- $[\![x]\!]_{\mathcal{M},\sigma} = \sigma(x)$.
- $[\![c]\!]_{\mathcal{M},\sigma} = c^{\mathcal{M}}$.
- $[\![f(t_1, ..., t_n)]\!]_{\mathcal{M},\sigma} = f^{\mathcal{M}}([\![t_1]\!]_{\mathcal{M},\sigma}, ..., [\![t_n]\!]_{\mathcal{M},\sigma})$.

The *satisfaction relation* $\mathcal{M} \models_\sigma \phi$ is given by the following clauses, this time assuming that the domain of σ contains all free variables of ϕ:

- $\mathcal{M} \models_\sigma t = t'$ iff $[\![t]\!]_{\mathcal{M},\sigma} = [\![t']\!]_{\mathcal{M},\sigma}$.
- $\mathcal{M} \models_\sigma R(t_1, ..., t_n)$ iff $R^{\mathcal{M}}([\![t_1]\!]_{\mathcal{M},\sigma}, ..., [\![t_n]\!]_{\mathcal{M},\sigma})$.
- $\mathcal{M} \models_\sigma \neg\phi$ iff $\mathcal{M} \not\models_\sigma \phi$.
- $\mathcal{M} \models_\sigma \phi \wedge \psi$ iff $\mathcal{M} \models_\sigma \phi$ and $\mathcal{M} \models_\sigma \psi$.
- $\mathcal{M} \models_\sigma \forall x.\phi$ iff for all u in the universe of $\mathcal{M}$ we have $\mathcal{M} \models_{\sigma, x \mapsto v} \phi$.

Note that if σ and σ' agree on the free variables of t, then $[\![t]\!]_{\mathcal{M},\sigma} = [\![t]\!]_{\mathcal{M},\sigma'}$. Likewise $\mathcal{M} \models_\sigma \phi$ if and only iff $\mathcal{M} \models_{\sigma'} \phi$, provided σ and σ' agree on the free variables of ϕ.

The *theory* of a model $\mathcal{M}$, written $\mathsf{Th}(\mathcal{M})$, is the set of all formulae made true by $\mathcal{M}$, i.e. $\mathsf{Th}(\mathcal{M}) = \{\phi \mid \mathcal{M} \models \phi\}$. We say two models $\mathcal{M}$ and $\mathcal{N}$ are *elementary equivalent* if $\mathsf{Th}(\mathcal{M}) = \mathsf{Th}(\mathcal{N})$. In first-order logic $\mathsf{Th}(\mathcal{M}) \subseteq \mathsf{Th}(\mathcal{N})$ already implies that $\mathcal{M}$ and $\mathcal{N}$ are elementary equivalent.

3 Cathoristic logic

In this section we introduce the syntax and semantics of cathoristic logic.

3.1 Syntax

Syntactically, cathoristic logic is a multi-modal logic with one new operator.

Definition 1 Let Σ be a non-empty set of *actions*. Actions are ranged over by $a, a', a_1, b, ...$, and A ranges over finite subsets of Σ. The *formulae* of cathoristic logic, ranged over by $\phi, \psi, \xi ...$, are given by the following grammar.

$$\phi \quad ::= \quad \top \mid \phi \wedge \psi \mid \langle a \rangle \phi \mid !A$$

The first three forms of ϕ are standard from Hennessy-Milner logic [17]: $\top$ is logical truth, $\wedge$ is conjunction, and $\langle a \rangle \phi$ means that the current state can transition via action a to a new state at which ϕ holds. Tantum A, written $!A$, is the key novelty of cathoristic logic. Asserting $!A$ means: in the current state at most the modalities $\langle a \rangle$ that satisfy $a \in A$ are permissible.

We assume that $\langle a \rangle \phi$ binds more tightly than conjunction, so $\langle a \rangle \phi \wedge \psi$ is short for $(\langle a \rangle \phi) \wedge \psi$. We often abbreviate $\langle a \rangle \top$ to $\langle a \rangle$. We define falsity $\bot$ as $!\emptyset \wedge \langle a \rangle$ where a is an arbitrary action in Σ. Hence, Σ must be non-empty. Note that, in the absence of negation, we cannot readily define disjunction, implication, or $[a]$ modalities by de Morgan duality.

Convention 1 *From now on we assume a fixed set Σ of actions, except where stated otherwise.*

3.2 Semantics

The semantics of cathoristic logic is close to Hennessy-Milner logic, but uses deterministic transition systems augmented with labels on states.

Definition 2 A *cathoristic transition system* is a triple $\mathcal{L} = (S, \rightarrow, \lambda)$, where $(S, \rightarrow)$ is a deterministic labelled transition system over Σ, and λ is a function from states to sets of actions (not necessarily finite), subject to the following constraints:

- For all states $s \in S$ it is the case that $\{a \mid \exists t \; s \xrightarrow{a} t\} \subseteq \lambda(s)$. We call this condition *admissibility*.
- For all states $s \in S$, $\lambda(s)$ is either finite or Σ. We call this condition *well-sizedness*.

The intended interpretation is that $\lambda(w)$ is the set of allowed actions emanating from w. The λ function is the semantic counterpart of the ! operator. The admissibility restriction is in place because transitions $s \xrightarrow{a} t$ where $a \notin$

$\lambda(s)$ would be saying that an a action is possible at s but at the same time prohibited at s. Well-sizedness is not a fundamental restriction but rather a convenient trick. Cathoristic transition systems have two kinds of states:

- States s without restrictions on outgoing transitions. Those are labelled with $\lambda(s) = \Sigma$.
- States s with restriction on outgoing transitions. Those are labelled by a finite set $\lambda(s)$ of actions.

Defining λ on all states and not just on those with restrictions makes some definitions and proofs slightly easier.

As with other modal logics, satisfaction of formulae is defined relative to a particular state in the transition system, giving rise to the following definition.

Definition 3 A *cathoristic model*, ranged over by $\mathfrak{M}, \mathfrak{M}', ...$, is a pair $(\mathcal{L}, s)$, where $\mathcal{L}$ is a cathoristic transition system $(S, \to, \lambda)$, and s is a state from S. We call s the *start state* of the model. An cathoristic model is a *tree* if the underlying transition system is a tree whose root is the start state.

Satisfaction of a formula is defined relative to a cathoristic model.

Definition 4 The *satisfaction relation* $\mathfrak{M} \models \phi$ is defined inductively by the following clauses, where we assume that $\mathfrak{M} = (\mathcal{L}, s)$ and $\mathcal{L} = (S, \to, \lambda)$.

$$
\begin{array}{lll}
\mathfrak{M} \models \top & & \\
\mathfrak{M} \models \phi \wedge \psi & \text{iff} & \mathfrak{M} \models \phi \text{ and } \mathfrak{M} \models \psi \\
\mathfrak{M} \models \langle a \rangle \phi & \text{iff} & \text{there is transition } s \xrightarrow{a} t \text{ such that } (\mathcal{L}, t) \models \phi \\
\mathfrak{M} \models !A & \text{iff} & \lambda(s) \subseteq A
\end{array}
$$

The first three clauses are standard. The last clause enforces the intended meaning of $!A$: the permissible modalities in the model are *at least as constrained* as required by $!A$. They may even be more constrained if the inclusion $\lambda(s) \subseteq A$ is proper. For infinite sets Σ of actions, allowing $\lambda(s)$ to return arbitrary infinite sets in addition to σ does not make a difference because A is finite by construction, so $\lambda(s) \subseteq A$ can never hold anyway for infinite $\lambda(s)$.

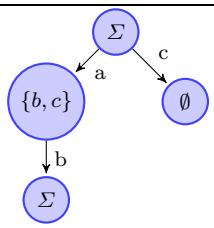

Fig. 1: Example model.

We continue with concrete examples. The model in Figure 1 satisfies all the following formulae, amongst others.

$\langle a\rangle$	$\langle a\rangle\langle b\rangle$	$\langle a\rangle!\{b,c\}$	$\langle a\rangle!\{b,c,d\}$	$\langle c\rangle$
$\langle c\rangle!\emptyset$	$\langle c\rangle!\{a\}$	$\langle c\rangle!\{a,b\}$	$\langle a\rangle \wedge \langle c\rangle$	$\langle a\rangle(\langle b\rangle\wedge!\{b,c\}$

Here we assume, as we do with all subsequent figures, that the top state is the start state. The same model does not satisfy any of the following formulae.

$$\langle b\rangle \qquad !\{a\} \qquad !\{a,c\} \qquad \langle a\rangle!\{b\} \qquad \langle a\rangle\langle c\rangle \qquad \langle a\rangle\langle b\rangle!\{c\}$$

Figure 2 shows various models of $\langle a\rangle\langle b\rangle$ and Figure 3 shows one model that does, and one that does not, satisfy the formula $!\{a,b\}$. Both models validate $!\{a,b,c\}$.

Cathoristic logic does not have the operators $\neg, \vee$, or $\rightarrow$. This has the following two significant consequences. First, every satisfiable formula has a unique (up to isomorphism) simplest model. In Figure 2, the left model is the unique simplest model satisfying$\langle a\rangle\langle b\rangle$. We will clarify below that model simplicity is closely related to the process theoretic concept of similarity, and use the existence of unique simplest models in our quadratic-time decision procedure.

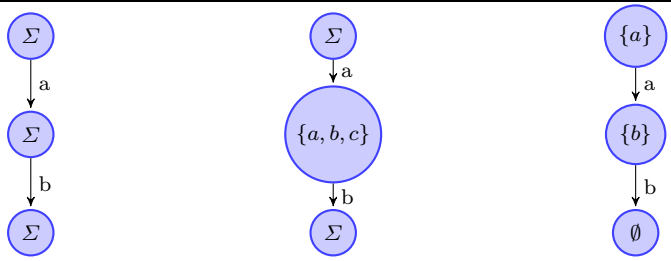

Fig. 2: Three models of $\langle a\rangle\langle b\rangle\top$

Fig. 3: The model on the left validates $!\{a,b\}$ while the model on the right does not.

Secondly, cathoristic logic is different from other logics in that there is an asymmetry between tautologies and contradictories: logics with conventional negation have an infinite number of non-trivial tautologies, as well as an infinite number of contradictories. In contrast, because cathoristic logic has no negation or disjunction operator, it is expressively limited in the tautologies it can express: $\top$ and conjunctions of $\top$ are its sole tautologies. On the other hand, the tantum operator enables an infinite number of contradictories to be expressed. For example:

$$\langle a\rangle \wedge !\emptyset \qquad \langle a\rangle \wedge !\{b\} \qquad \langle a\rangle \wedge !\{b,c\} \qquad \langle b\rangle \wedge !\emptyset$$

Next, we present the semantic consequence relation.

Definition 5 We say the formula ϕ *semantically implies* ψ, written $\phi \models \psi$, provided for all cathoristic models $\mathfrak{M}$ if it is the case that $\mathfrak{M} \models \phi$ then also $\mathfrak{M} \models \phi$. We sometimes write $\models \phi$ as a shorthand for $\top \models \phi$.

Cathoristic logic shares with other (multi)-modal logics the following implications:

$$\langle a\rangle\langle b\rangle \models \langle a\rangle \qquad\qquad \langle a\rangle(\langle b\rangle \wedge \langle c\rangle) \models \langle a\rangle\langle b\rangle$$

As cathoristic logic is restricted to deterministic models, it also validates the following formula:

$$\langle a\rangle\langle b\rangle \wedge \langle a\rangle\langle b\rangle \models \langle a\rangle(\langle b\rangle \wedge \langle c\rangle)$$

Cathoristic logic also validates all implications in which the set of constraints is relaxed from left to right. For example:

$$!\{c\} \models !\{a,b,c\} \qquad\qquad !\emptyset \models !\{a,b\}$$

4 Inferences between atomic sentences

Cathoristic logic arose in part as an attempt to answer the question: what is the simplest logic that can capture inferences between atomic sentences of natural language? In this section, we give examples of such inferences, and then show how cathoristic logic handles them. We also compare our approach with attempts at expressing the inferences in first-order logic.

4.1 Intra-atomic inferences in cathoristic logic

Natural language admits many types of inference between atomic sentences. First, exclusion:

> "Jack is male" is incompatible with "Jack is female".

Second, entailment inferences from dyadic to monadic predicates:

> "Jack loves Jill" implies "Jack loves".

Third, adverbial inferences:

> "Jack walks quickly" implies "Jack walks".

Fourth, inferences from conjunctions of sentences to conjunctions of noun-phrases (and vice-versa):

> "Jack loves Jill" and "Jack loves Joan" together imply that "Jack loves Jill and Joan".

Fifth, inferences from conjunctions of sentences to conjunction of predicates[10] (and vice-versa):

> "Jack is bruised" and "Jack is humiliated" together imply that "Jack is bruised and humiliated".

They all can be handled directly and naturally in cathoristic logic, as we shall now show.

Incompatibility, such as that between "Jack is male" and "Jack is female", is translated into cathoristic logic as the pair of incompatible sentences:

$$\langle jack\rangle\langle sex\rangle(\langle male\rangle\wedge !\{male\}) \qquad \langle jack\rangle\langle sex\rangle(\langle female\rangle\wedge !\{female\}).$$

[10] See [28] p.282 for a spirited defence of predicate conjunction against Fregean regimentation.

Cathoristic logic handles entailments from dyadic to monadic predicates[11]. "Jack loves Jill" is translated into cathoristic logic as:

$$\langle jack\rangle\langle loves\rangle\langle jill\rangle.$$

The semantics of modalities ensures that this directly entails:

$$\langle jack\rangle\langle loves\rangle.$$

Similarly, cathoristic logic supports inferences from triadic to dyadic predicates:

"Jack passed the biscuit to Mary" implies "Jack passed the biscuit".

This can be expressed directly in cathoristic logic as:

$$\langle jack\rangle\langle passed\rangle\langle biscuit\rangle\langle to\rangle(\langle mary\rangle\wedge!\{mary\}) \models \langle jack\rangle\langle passed\rangle\langle biscuit\rangle.$$

Adverbial inferences is captured in cathoristic logic as follows.

$$\langle jack\rangle\langle walks\rangle\langle quickly\rangle$$

entails:

$$\langle jack\rangle\langle walks\rangle.$$

Cathoristic logic directly supports inferences from conjunctions of sentences to conjunctions of noun-phrases. As our models are deterministic, we have the general rule that $\langle a\rangle\langle b\rangle \wedge \langle a\rangle\langle c\rangle \models \langle a\rangle(\langle b\rangle \wedge \langle c\rangle)$ from which it follows that

$$\langle jack\rangle\langle loves\rangle\langle jill\rangle \qquad \text{and} \qquad \langle jack\rangle\langle loves\rangle\langle joan\rangle$$

together imply

$$\langle jack\rangle\langle loves\rangle(\langle jill\rangle \wedge \langle joan\rangle).$$

Using the same rule, we can infer that

$$\langle jack\rangle\langle bruised\rangle \wedge \langle jack\rangle\langle humiliated\rangle$$

together imply

$$\langle jack\rangle(\langle bruised\rangle \wedge \langle humiliated\rangle).$$

[11] Although natural languages are full of examples of inferences from dyadic to monadic predicates, there are certain supposed counterexamples to the general rule that a dyadic predicate always implies a monadic one. For example, "Jack explodes the device" does not, on its most natural reading, imply that "Jack explodes". Our response to cases like this is to distinguish between two distinct monadic predicates $explodes_1$ and $explodes_2$:

- $X explodes_1$ iff X is an object that undergoes an explosion
- $X explodes_2$ iff X is an agent that initiates an explosion

Now "Jack explodes the device" does imply that "Jack $explodes_2$" but does not imply that "Jack $explodes_1$". There is no deep problem here - just another case where natural language overloads the same word in different situation to have different meanings.

4.2 Intra-atomic inferences in first-order logic

Next, we look at how these inferences are handled in first-order logic.

4.2.1 Incompatible predicates in first-order logic

How are incompatible predicates represented in first-order logic? Brachman and Levesque [6] introduce the topic by remarking:

> We would consider it quite "obvious" in this domain that if it were asserted that $John$ were a Man, then we should answer "no" to the query $Woman(John)$.

They propose adding an extra axiom to express the incompatibility:

$$\forall x.(Man(x) \to \neg Woman(x))$$

This proposal imposes a burden on the knowledge-representer: an extra axiom must be added for every pair of incompatible predicates. This is burdensome for large sets of incompatible predicates. For example, suppose there are 50 football teams, and a person can only support one team at a time. We would need to add $\binom{50}{2}$ axioms, which is unwieldy.

$$\begin{array}{c}\forall x.\neg(SupportsArsenal(x) \wedge SupportsLiverpool(x))\\ \forall x.\neg(SupportsArsenal(x) \wedge SupportsManUtd(x))\\ \forall x.\neg(SupportsLiverpool(x) \wedge SupportsManUtd(x))\\ \vdots\end{array}$$

Or, if we treat the football-teams as objects, and have a two-place $Supports$ relation between people and teams, we could have:

$$\forall xyz.(Supports(x,y) \wedge y \neq z \to \neg Supports(x,z)).$$

If we also assume that each football team is distinct from all others, this certainly captures the desired uniqueness condition. But it does so by using relatively complex logical machinery.

4.2.2 Inferences from dyadic to monadic predicates in first-order logic

If we want to capture the inference from "Jack loves Jill" to "Jack loves" in first-order logic, we can use a non-logical axiom:

$$\forall x.y.(Loves_2(x,y) \to Loves_1(x))$$

We would have to add an extra axiom like this for every n-place predicate. This is cumbersome at best. In cathoristic logic, by contrast, we do not need to introduce any non-logical machinery to capture these inferences because they all follow from the general rule that $\langle a\rangle\langle b\rangle \models \langle a\rangle$.

4.2.3 Adverbial inferences in first-order logic

How can we represent verbs in traditional first-order logic so as to support adverbial inference? Davidson [10] proposes that every n-place action verb be analysed as an n+1-place predicate, with an additional slot representing an event. For example, he analyses "I flew my spaceship to the Morning Star" as

$$\exists x.(Flew(I, MySpaceship, x) \wedge To(x, TheMorningStar))$$

This implies

$$\exists x.Flew(I, MySpaceship, x)$$

This captures the inference from "I flew my spaceship to the Morning Star" to "I flew my spaceship".

First-order logic cannot support logical inferences between atomic sentences. If it is going to support inferences from adverbial sentences, it *cannot* treat them as atomic and must instead *reinterpret* them as logically complex propositions. The cost of Davidson's proposal is that a seemingly simple sentence - such as "Jones walks" - turns out, on closer inspection, not to be atomic at all - but to involve existential quantification:

$$\exists x.Walks(Jones, x)$$

First-order logic *can* handle such inferences - but only by reinterpreting the sentences as logically-complex compound propositions.

5 Cathoristic logic as a language for knowledge representation

Cathoristic logic has been used as the representation language for a large, complex, dynamic multi-agent simulation [13]. This is an industrial-sized application, involving tens of thousands of rules and facts[12]. In this simulation, the entire world state is stored as a cathoristic model.

We found that cathoristic logic has two distinct advantages as a language for knowledge representation. First, it is ergonomic: ubiquitous concepts (such as uniqueness) can be expressed directly. Second, it is efficient: the tantum operator allows certain sorts of optimisation that would not otherwise be available. We shall consider these in turn.

5.1 Representing facts in cathoristic logic

A sentence involving a one-place predicate of the form $p(a)$ is expressed in cathoristic logic as

$$\langle a\rangle\langle p\rangle$$

A sentence involving a many-to-many two-place relation of the form $r(a, b)$ is expressed in cathoristic logic as

$$\langle a\rangle\langle r\rangle\langle b\rangle$$

But a sentence involving a many-to-one two-place relation of the form $r(a, b)$ is expressed as:

$$\langle a\rangle\langle r\rangle(\langle b\rangle\wedge!\{b\})$$

So, for example, to say that "Jack likes Jill" (where "likes" is, of course, a many-many relation), we would write:

$$\langle jack\rangle\langle likes\rangle\langle jill\rangle$$

But to say that "Jack is married to Joan" (where"is-married-to" is a many-one relation), we would write:

$$\langle jack\rangle\langle married\rangle(\langle joan\rangle\wedge!\{joan\})$$

Colloquially, we might say that "Jack is married to Joan - and only Joan". Note that the relations are placed in infix position, so that the facts about an object are "contained" within the object. One reason for this particular way of structuring the data will be explained below.

Consider the following facts about a gentleman named Brown:

[12] The application has thousands of paying users, and available for download on the App Store for the iPad [12].

$$\langle brown\rangle \left(\begin{array}{l} \langle sex\rangle(\langle male\rangle \wedge !\{male\}) \\ \qquad \wedge \\ \langle friends\rangle(\langle lucy\rangle \wedge \langle elizabeth\rangle) \end{array} \right)$$

All facts starting with the prefix $\langle brown\rangle$ form a sub-tree of the entire database. And all facts which start with the prefix $\langle brown\rangle\langle friends\rangle$ form a sub-tree of that tree. A sub-tree can be treated as an individual via its prefix. A sub-tree of formulae is the cathoristic logic equivalent of an *object* in an object-oriented programming language.

To model change over time, we assert and retract statements from the database, using a non-monotonic update mechanism. If a fact is inserted into the database that involves a state-labelling restricting the permissible transitions emanating from that state, then all transitions out of that state that are incompatible with the restriction are removed. So, for example, if the database currently contains the fact that the traffic light is amber, and then we update the database to assert the traffic light is red:

$$\langle tl\rangle\langle colour\rangle(\langle red\rangle \wedge !\{red\})$$

Now the restriction on the state (that red is the only transition) means that the previous transition from that state (the transition labelled with amber) is automatically removed.

The tree-structure of formulae allows us to express the *life-time of data* in a natural way. If we wish a piece of data d to exist for just the duration of a proposition t, then we make t be a sub-expression of d. For example, if we want the friendships of an agent to exist just as long as the agent, then we place the relationships inside the agent:

$$\langle brown\rangle\langle friends\rangle$$

Now, when we remove $\langle brown\rangle$ all the sub-trees, including the data about who he is friends with, will be automatically deleted as well.

Another advantage of our representation is that we get a form of *automatic currying* which simplifies queries. So if, for example, Brown is married to Elizabeth, then the database would contain

$$\langle brown\rangle\langle married\rangle(\langle elizabeth\rangle \wedge !\{elizabeth\})$$

In cathoristic logic, if we want to find out whether Brown is married, we can query the sub-formula directly - we just ask if

$$\langle brown\rangle\langle married\rangle$$

In first-order logic, if *married* is a two-place predicate, then we need to fill in the extra argument place with a free variable - we would need to find out if there exists an x such that $married(brown, x)$ - this is more cumbersome to type and slower to compute.

5.2 Simpler postconditions

In this section, we contrast the representation in action languages based on first-order logic[13], with our cathoristic logic-based representation. Action definitions are rendered in typewriter font.

When expressing the pre- and postconditions of an action, planners based on first-order logic have to explicitly describe the propositions that are removed when an action is performed:

```
action move(A, X, Y)
    preconditions
        at(A, X)
    postconditions
        add: at(A, Y)
        remove: at(A, X)
```

Here, we need to explicitly state that when A moves from X to Y, A is no longer at X. It might seem obvious to us that if A is now at Y, he is no longer at X - but we need to explicitly tell the system this. This is unnecessary, cumbersome and error-prone. In cathoristic logic, by contrast, the exclusion operator means we do not need to specify the facts that are no longer true:

```
action move (A, X, Y)
    preconditions
        <A><at>(<X> /\ !{X})
    postconditions
        add: <A><at>(<Y> /\ !{Y})
```

The tantum operator ! makes it clear that something can only be at one place at a time, and the non-monotonic update rule described above *automatically* removes the old invalid location data.

5.3 Using tantum ! to optimise preconditions

Suppose, for example, we want to find all married couples who are both Welsh. In Prolog, we might write something like:

```
welsh_married_couple(X, Y) :-
    welsh(X),
    welsh(Y),
    spouse(X,Y).
```

Rules like this create a large search-space because we need to find all instances of $welsh(X)$ and all instances of $welsh(Y)$ and take the cross-product [27]. If there are n Welsh people, then we will be searching n^2 instances of (X, Y) substitutions.

[13] E.g. STRIPS [14]

If we express the rule in cathoristic logic, the compiler is able to use the extra information expressed in the ! operator to reorder the literals to find the result significantly faster. Assuming someone can only have a single spouse at any moment, the rule is expressed in cathoristic logic as:

```
welsh_married_couple(X, Y) :-
    <welsh> <X>,
    <welsh> <Y>,
    <spouse> <X> (<Y> /\ !{Y}).
```

Now the compiler is able to reorder these literals to minimise the search-space. It can see that, once X is instantiated, the following literal can be instantiated without increasing the search-space:

```
<spouse> <X> (<Y> /\ !{Y})
```

The *tantum* operator can be used by the compiler to see that there is at most one Y who is the spouse of X. So the compiler reorders the clauses to produce:

```
welsh_married_couple (X, Y) :-
    <welsh> <X>,
    <spouse> <X> (<Y> /\ !{Y}),
    <welsh> <Y>.
```

Now it is able to find all results by just searching n instances - a significant optimisation. In our application, this optimisation has made a significant difference to the run-time cost of query evaluation.

6 Semantics and Decision Procedure

In this section we provide our key semantic results. We define a partial ordering $\preceq$ on models, and show how the partial ordering can be extended into a bounded lattice. We use the bounded lattice to construct a quadratic-time decision procedure.

6.1 Semantic characterisation of elementary equivalence

Elementary equivalence induces a notion of model equivalence: two models are elementarily equivalent exactly when they make the same formulae true. Elementary equivalence as a concept thus relies on cathoristic logic even for its definition. We now present an alternative characterisation that is purely semantic, using the concept of (mutual) simulation from process theory. Apart from its intrinsic interest, this characterisation will also be crucial for proving completeness of the proof rules.

We first define a pre-order $\preceq$ on models by extending the notion of simulation on labelled transition systems to cathoristic models. Then we prove an alternative characterisation of $\preceq$ in terms of set-inclusion of the theories induced by models. We then show that two models are elementarily equivalent exactly when they are related by $\preceq$ and by $\preceq^{-1}$.

Definition 6 Let $\mathcal{L}_i = (S_i, \rightarrow_i, \lambda_i)$ be cathoristic transition systems for $i = 1, 2$. A relation $\mathcal{R} \subseteq S_1 \times S_2$ is a *simulation from* $\mathcal{L}_1$ *to* $\mathcal{L}_2$, provided:

- $\mathcal{R}$ is a simulation on the underlying transition systems.
- Whenever $(x, y) \in \mathcal{R}$ then also $\lambda_1(x) \supseteq \lambda_2(y)$.

If $\mathfrak{M}_i = (\mathcal{L}_i, x_i)$ are models, we say $\mathcal{R}$ is a *simulation from* $\mathfrak{M}_1$ *to* $\mathfrak{M}_2$, provided the following hold.

- $\mathcal{R}$ is a simulation from $\mathcal{L}_1$ to $\mathcal{L}_2$ as cathoristic transition systems.
- $(x_1, x_2) \in \mathcal{R}$.

Note that the only difference from the usual definition of simulation is the additional requirement on the state labelling functions λ_1 and λ_2.

Definition 7 The largest simulation from $\mathfrak{M}_1$ to $\mathfrak{M}_2$ is denoted $\mathfrak{M}_1 \preceq_{sim} \mathfrak{M}_2$. It is easy to see that $\preceq_{sim}$ is itself a simulation from $\mathfrak{M}_1$ to $\mathfrak{M}_2$, and the union of all such simulations. If $\mathfrak{M}_1 \preceq_{sim} \mathfrak{M}_2$ we say $\mathfrak{M}_2$ *simulates* $\mathfrak{M}_1$.

We write $\simeq$ for $\preceq_{sim} \cap \preceq_{sim}^{-1}$. We call $\simeq$ the *mutual simulation* relation.

We briefly discuss the relationship of $\simeq$ with bisimilarity, a notion of equality well-known from process theory and modal logic. For non-deterministic transition systems $\simeq$ is a strictly coarser relation than bisimilarity.

Definition 8 We say $\mathcal{R}$ is a *bisimulation* if $\mathcal{R}$ is a simulation from $\mathfrak{M}_1$ to $\mathfrak{M}_2$ and $\mathcal{R}^{-1}$ is a simulation from $\mathfrak{M}_2$ to $\mathfrak{M}_1$. By $\sim$ we denote the largest bisimulation, and we say that $\mathfrak{M}_1$ and $\mathfrak{M}_2$ are *bisimilar* whenever $\mathfrak{M}_1 \sim \mathfrak{M}_2$.

Lemma 1 *On cathoristic models,* $\sim$ *and* $\simeq$ *coincide.*

Proof: Straightforward from the definitions. □

Definition 9 Let $\mathsf{Th}(\mathfrak{M})$ be the *theory* of $\mathfrak{M}$, i.e. the formulae made true by $\mathfrak{M}$, i.e. $\mathsf{Th}(\mathfrak{M}) = \{\phi \mid \mathfrak{M} \models \phi\}$.

We give an alternative characterisation on $\preceq_{sim}^{-1}$ using theories. In what follows, we will mostly be interested in $\preceq_{sim}^{-1}$, so we give it its own symbol.

Definition 10 Let $\preceq$ be short for $\preceq_{sim}^{-1}$.

Figure 4 gives some examples of models and how they are related by $\preceq$.

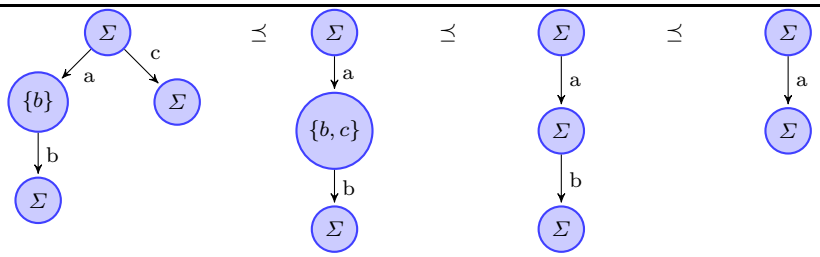

Fig. 4: Examples of $\preceq$

Theorem 1 (Characterisation of elementary equivalence)

1. $\mathfrak{M}' \preceq \mathfrak{M}$ *if and only if* $\mathsf{Th}(\mathfrak{M}) \subseteq \mathsf{Th}(\mathfrak{M}')$.
2. $\mathfrak{M}' \simeq \mathfrak{M}$ *if and only if* $\mathsf{Th}(\mathfrak{M}) = \mathsf{Th}(\mathfrak{M}')$.

Proof: For (1) assume $\mathfrak{M}' \preceq \mathfrak{M}$ and $\mathfrak{M} \models \phi$. We must show $\mathfrak{M}' \models \phi$. Let $\mathfrak{M} = (\mathcal{L}, w)$ and $\mathfrak{M}' = (\mathcal{L}', w')$. The proof proceeds by induction on ϕ. The cases for $\top$ and $\wedge$ are trivial. Assume $\phi = \langle a \rangle \psi$ and assume $(\mathcal{L}, w) \models \langle a \rangle \psi$. Then $w \xrightarrow{a} x$ and $(\mathcal{L}, x) \models \psi$. As $\mathfrak{M}'$ simulates $\mathfrak{M}$, there is an x' such that $(x, x') \in R$ and $w' \xrightarrow{a} x'$. By the induction hypothesis, $(\mathcal{L}', x') \models \psi$. Therefore, by the semantic clause for $\langle\rangle$, $(\mathcal{L}', w') \models \langle a \rangle \psi$. Assume now that $\phi = !\ A$, for some finite $A \subseteq \Sigma$, and that $(\mathcal{L}, w) \models !\ A$. By the semantic clause for !, $\lambda(w) \subseteq A$. Since $(\mathcal{L}', w') \preceq (\mathcal{L}, w)$, by the definition of simulation of cathoristic transition systems, $\lambda(w) \supseteq \lambda'(w')$. Therefore, $\lambda'(w') \subseteq \lambda(w) \subseteq A$. Therefore, by the semantic clause for !, $(\mathcal{L}', w') \models !\ A$.

For the other direction, let $\mathfrak{M} = (\mathcal{L}, w)$ and $\mathfrak{M}' = (\mathcal{L}', w')$. Assume $\mathsf{Th}(\mathfrak{M}) \subseteq \mathsf{Th}(\mathfrak{M}')$. We need to show that $\mathfrak{M}'$ simulates $\mathfrak{M}$. In other words, we need to produce a relation $R \subseteq S \times S'$ where S is the state set of $\mathcal{L}$, S' is the state set for $\mathcal{L}'$ and $(w, w') \in R$ and R is a simulation from $(\mathcal{L}, w)$ to $(\mathcal{L}', w')$. Define $R = \{(x, x') \mid \mathsf{Th}((\mathcal{L}, x)) \subseteq \mathsf{Th}((\mathcal{L}', x'))\}$. Clearly, $(w, w') \in R$, as $\mathsf{Th}((\mathcal{L}, w)) \subseteq \mathsf{Th}((\mathcal{L}', w'))$. To show that R is a simulation, assume $x \xrightarrow{a} y$ in $\mathcal{L}$ and $(x, x') \in R$. We need to provide a y' such that $x' \xrightarrow{a} y'$ in $\mathcal{L}'$ and

$(y, y') \in R$. Consider the formula $\langle a \rangle \mathsf{char}((\mathcal{L}, y))$. Now $x \models \langle a \rangle \mathsf{char}((\mathcal{L}, y))$, and since $(x, x') \in R$, $x' \models \langle a \rangle \mathsf{char}((\mathcal{L}, y))$. By the semantic clause for $\langle a \rangle$, if $x' \models \langle a \rangle \mathsf{char}((\mathcal{L}, y))$ then there is a y' such that $y' \models \mathsf{char}((\mathcal{L}, y))$. We need to show $(y, y') \in R$, i.e. that $y \models \phi$ implies $y' \models \phi$ for all ϕ. Assume $y \models \phi$. Then by the definition of $\mathsf{char}()$, $\mathsf{char}((\mathcal{L}, y)) \models \phi$. Since $y' \models \mathsf{char}((\mathcal{L}, y))$, $y' \models \phi$. So $(y, y') \in R$, as required.

Finally,we need to show that whenever $(x, x') \in R$, then $\lambda(x) \supseteq \lambda'(x')$. Assume, first, that $\lambda(x)$ is finite. Then $(\mathcal{L}, x) \models ! \ \lambda(x)$. But as $(x, x') \in R$, $\mathsf{Th}((\mathcal{L}, x)) \subseteq \mathsf{Th}((\mathcal{L}', x'))$, so $(\mathcal{L}', x') \models ! \ \lambda(x)$. But, by the semantic clause for !, $(\mathcal{L}', x') \models ! \ \lambda(x)$ iff $\lambda'(x') \subseteq \lambda(x)$. Therefore $\lambda(x) \supseteq \lambda'(x')$. If, on the other hand, $\lambda(x)$ is infinite, then $\lambda(x) = \Sigma$ (because the only infinite state labelling that we allow is Σ). Every state labelling is a subset of Σ, so here too, $\lambda(x) = \Sigma \supseteq \lambda'(x')$.

This establishes (1), and (2) is immediate from the definitions.

□

Theorem 1.1 captures one way in which the model theory of classical and cathoristic logic differ. In classical logic the theory of each model is complete, and $\mathsf{Th}(\mathcal{M}) \subseteq \mathsf{Th}(\mathcal{N})$ already implies that $\mathsf{Th}(\mathcal{M}) = \mathsf{Th}(\mathcal{N})$, i.e. $\mathcal{M}$ and $\mathcal{N}$ are elementarily equivalent. Cathoristic logic's lack of negation changes this drastically, and gives $\preceq$ the structure of a non-trivial bounded lattice as we shall demonstrate below.

Theorem 1 has various consequences.

Corollary 1 *1. If ϕ has a model then it has a model whose underlying transition system is a tree, i.e. all states except for the start state have exactly one predecessor, and the start state has no predecessors.*
2. If ϕ has a model then it has a model where every state is reachable from the start state.

Proof: Both are straightforward because $\simeq$ is closed under tree-unfolding as well as under removal of states not reachable from the start state. □

6.2 Quotienting models

The relation $\preceq$ is not a partial order, only a pre-order. For example

$$\mathfrak{M}_1 = (\{w\}, \emptyset, \{w \mapsto \Sigma\}), w) \qquad \mathfrak{M}_2 = (\{v\}, \emptyset, \{v \mapsto \Sigma\}), v)$$

are two distinct models with $\mathfrak{M}_1 \preceq \mathfrak{M}_2$ and $\mathfrak{M}_2 \preceq \mathfrak{M}_1$. The difference between the two models, the name of the unique state, is trivial and not relevant for the formulae they make true: $\mathsf{Th}(\mathfrak{M}_1) = \mathsf{Th}(\mathfrak{M}_2)$. As briefly mentioned in the mathematical preliminaries (Section 2), we obtain a proper partial-order by simply quotienting models:

$$\mathfrak{M} \simeq \mathfrak{M}' \qquad \text{iff} \qquad \mathfrak{M} \preceq \mathfrak{M}' \text{ and } \mathfrak{M}' \preceq \mathfrak{M}$$

and then ordering the $\simeq$-equivalence classes as follows:

$$[\mathfrak{M}]_\simeq \preceq [\mathfrak{M}']_\simeq \qquad \text{iff} \qquad \mathfrak{M} \preceq \mathfrak{M}'.$$

Greatest lower and least upper bounds can also be computed on representatives:

$$\bigsqcup\{[\mathfrak{M}]_\simeq \mid \mathfrak{M} \in S\ \} \quad = \quad [\bigsqcup S]_\simeq$$

whenever $\bigsqcup S$ exists, and likewise for the greatest lower bound. We also define

$$[\mathfrak{M}]_\simeq \models \phi \qquad \text{iff} \qquad \mathfrak{M} \models \phi.$$

It is easy to see that these definitions are independent of the chosen representatives.

In the rest of this text we will usually be sloppy and work with concrete models instead of $\simeq$-equivalence classes of models because the quotienting process is straightforward and not especially interesting. We can do this because all relevant subsequent constructions are also representation independent.

6.3 The bounded lattice of models

It turns out that $\preceq$ on ($\simeq$-equivalence classes of) models is not just a partial order, but a bounded lattice, except that a bottom element is missing.

Definition 11 We extend the collection of models with a single *bottom* element $\bot$, where $\bot \models \phi$ for all ϕ. We also write $\bot$ for $[\bot]_\simeq$. We extend the relation $\preceq$ and stipulate that $\bot \preceq \mathfrak{M}$ for all models $\mathfrak{M}$.

Theorem 2 *The collection of (equivalence classes of) models together with $\bot$, and ordered by $\preceq$ is a bounded lattice.*

Proof: The topmost element in the lattice is the model $((\{w\}, \emptyset, \{w \mapsto \Sigma\}), w)$ (for some state w): this is the model with no transitions and no transition restrictions. The bottom element is $\bot$. Below, we shall define two functions glb and lub, and show that they satisfy the required properties of $\sqcap$ and $\sqcup$ respectively. □

Cathoristic logic is unusual in that every set of models has a unique (up to isomorphism) least upper bound. Logics with disjunction, negation or implication do not have this property.

Consider propositional logic, for example. Define a model of propositional logic as a set of atomic formulae that are set to true. Then we have a natural ordering on propositional logic models:

$$\mathfrak{M} \leq \mathfrak{M}' \quad \text{iff} \quad \mathfrak{M} \supseteq \mathfrak{M}'$$

Consider all the possible models that satisfy $\phi \vee \psi$:

$$\{\phi\} \qquad \{\psi\} \qquad \{\phi, \psi\} \qquad \{\phi, \psi, \xi\} \qquad \cdots$$

This set of satisfying models has no least upper bound, since $\{\phi\} \not\leq \{\psi\}$ and $\{\psi\} \not\leq \{\phi\}$. Similarly, the set of models satisfying $\neg(\neg\phi \wedge \neg\psi)$ has no least upper bound.

The fact that cathoristic logic models have unique least upper bounds is used in proving completeness of our inference rules, and implementing the quadratic-time decision procedure.

6.4 Computing the least upper bound of the models that satisfy a formula

In our decision procedure, we will see if $\phi \models \psi$ by constructing the least upper bound of the models satisfying ϕ, and checking whether it satisfies ψ.

In this section, we define a function $\mathsf{simpl}(\phi)$ that satisfies the following condition:

$$\mathsf{simpl}(\phi) = \bigsqcup\{\mathfrak{M} | \mathfrak{M} \models \phi\}$$

Define $\mathsf{simpl}(\phi)$ as:

$$\begin{aligned}
\mathsf{simpl}(\top) &= ((\{v\}, \emptyset, \{v \mapsto \Sigma\}), v) \\
\mathsf{simpl}(!A) &= ((\{v\}, \emptyset, \{v \mapsto A\}), v) \\
\mathsf{simpl}(\phi_1 \wedge \phi_2) &= \mathsf{glb}(\mathsf{simpl}(\phi_1), \mathsf{simpl}(\phi_2)) \\
\mathsf{simpl}(\langle a \rangle \phi) &= ((S \cup \{w'\}, \rightarrow \cup (w' \xrightarrow{a} w), \lambda \cup \{w' \mapsto \Sigma\}]), w') \\
&\quad \text{where } \mathsf{simpl}(\phi) = ((S, \rightarrow, \lambda), w) \text{and } w' \text{ is a new state} \\
&\quad \text{not appearing in } S
\end{aligned}$$

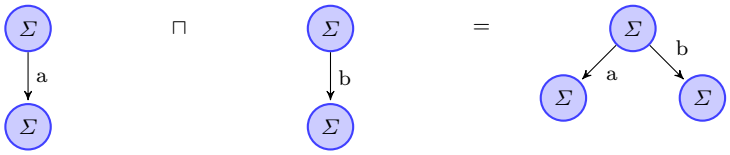

Fig. 5: Example of ⊓.

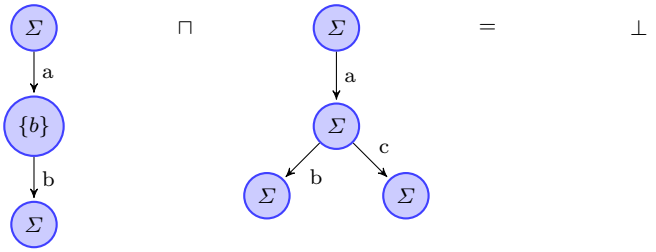

Fig. 6: Example of ⊓.

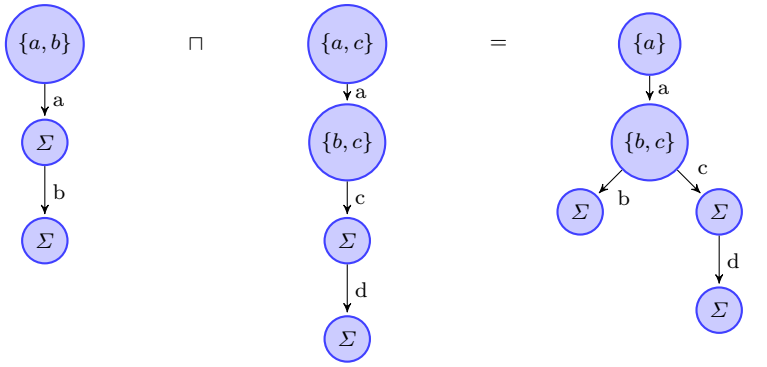

Fig. 7: Example of ⊓.

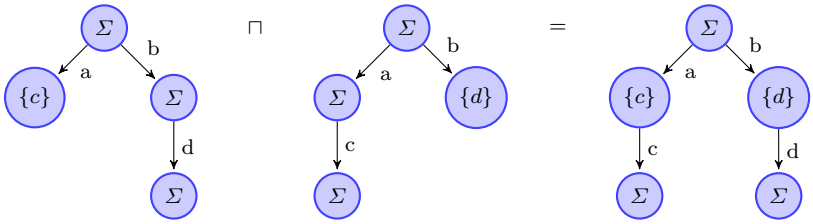

Fig. 8: Example of ⊓.

Note that, by our conventions, $\mathsf{simpl}(\phi)$ really returns a $\simeq$-equivalence class of models.

The only complex case is the clause for $\mathsf{simpl}(\phi_1 \wedge \phi_2)$, which uses the glb function, defined as follows, where we assume that the sets of states in the two models are disjoint and are trees. It is easy to see that $\mathsf{simpl}(\cdot)$ always returns tree models.

$$\begin{aligned}
\mathsf{glb}(\bot, \mathfrak{M}) &= \bot \\
\mathsf{glb}(\mathfrak{M}, \bot) &= \bot \\
\mathsf{glb}(\mathfrak{M}, \mathfrak{M}') &= \mathsf{merge}(\mathcal{L}, \mathcal{L}', \{(w, w')\}) \\
&\quad \text{where } \mathfrak{M} = (\mathcal{L}, w) \text{ and } \mathfrak{M}' = (\mathcal{L}', w')
\end{aligned}$$

The merge function returns $\bot$ if either of its arguments are $\bot$. Otherwise, it merges the two transition systems together, given a set of state-identification pairs (a set of pairs of states from the two transition systems that need to be identified). The state-identification pairs are used to make sure that the resulting model is deterministic.

$$\mathsf{merge}(\mathcal{L},\mathcal{L}',ids) = \begin{cases} \bot & \text{if } \mathsf{inconsistent}(\mathcal{L},\mathcal{L}',ids) \\ \mathsf{join}(\mathcal{L},\mathcal{L}') & \text{if } ids = \emptyset \\ \mathsf{merge}(\mathcal{L},\mathcal{L}'',ids') & \text{else, where } \mathcal{L}'' = \mathsf{applyIds}(ids,\mathcal{L}') \\ & \text{and } ids' = \mathsf{getIds}(\mathcal{L},\mathcal{L}',ids) \end{cases}$$

The inconsistent predicate is true if there is pair of states in the state-identification set such that the out-transitions of one state is incompatible with the state-labelling on the other state:

$\mathsf{inconsistent}(\mathcal{L},\mathcal{L}',ids)$
iff $\exists (w,w') \in ids$ with $\mathsf{out}(\mathcal{L},w) \not\subseteq \lambda'(w')$ or $\mathsf{out}(\mathcal{L}',w') \not\subseteq \lambda(w)$.

Here the out function returns all the actions immediately available from the given state w.

$$\mathsf{out}(((S,\to,\lambda),w)) \;=\; \{a \mid \exists w'.w \xrightarrow{a} w'\}$$

The join function takes the union of the two transition systems.

$$\mathsf{join}((S,\to,\lambda),(S',\to',\lambda')) \quad = \quad (S \cup S', \to \cup \to', \lambda'')$$

Here λ'' takes the constraints arising from both, λ and λ' into account:

$$\begin{array}{rcl} & & \{\lambda(s) \cap \lambda'(s) \mid s \in S \cup S'\} \\ \lambda''(s) & = & \cup\, \{\lambda(s) \mid s \in S \setminus S'\} \\ & & \cup\, \{\lambda(s) \mid s \in S' \setminus S\}. \end{array}$$

The applyIds function applies all the state-identification pairs as substitutions to the Labelled Transition System:

$$\mathsf{applyIds}(ids,(S,\to,\lambda)) \quad = \quad (S',\to',\lambda')$$

where

$$\begin{array}{rcl} S' & = & S\,[w/w' \mid (w,w') \in ids] \\ \to' & = & \to\,[w/w' \mid (w,w') \in ids] \\ \lambda' & = & \lambda\,[w/w' \mid (w,w') \in ids] \end{array}$$

Here $[w/w' \mid (w,w') \in ids]$ means the simultaneous substitution of w for w' for all pairs (w,w') in ids. The getIds function returns the set of extra state-identification pairs that need to be added to respect determinism:

$$\mathsf{getIds}(\mathcal{L},\mathcal{L}',ids) \quad = \quad \{(x,x') \mid (w,w') \in ids, \exists a\,.\, w \xrightarrow{a} x, w' \xrightarrow{a} x'\}$$

The function $\mathsf{simpl}(\cdot)$ has the expected properties, as the next lemma shows.

Lemma 2 $\mathsf{simpl}(\phi) \models \phi$.

Proof: By induction on ϕ. □

Lemma 3 glb *as defined is the greatest lower bound*

We will show that:

- $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}') \preceq \mathfrak{M}$ and $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}') \preceq \mathfrak{M}'$
- If $\mathfrak{N} \preceq \mathfrak{M}$ and $\mathfrak{N} \preceq \mathfrak{M}'$, then $\mathfrak{N} \preceq \mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$

If $\mathfrak{M}$, $\mathfrak{M}'$ or $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$ are equal to $\bot$, then we just apply the rule that $\bot \preceq m$ for all models m. So let us assume that $\mathsf{consistent}(\mathfrak{M}, \mathfrak{M}')$ and that $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}') \neq \bot$.

Proof: To show $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}') \preceq \mathfrak{M}$, we need to provide a simulation $\mathcal{R}$ from $\mathfrak{M}$ to $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$. If $\mathfrak{M} = ((S, \rightarrow, \lambda), w)$, then define $\mathcal{R}$ as the identity relation on the states of S:

$$\mathcal{R} = \{(x, x) \mid x \in S\}$$

It is straightforward to show that $\mathcal{R}$ as defined is a simulation from $\mathfrak{M}$ to $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$. If there is a transition $x \xrightarrow{a} y$ in $\mathfrak{M}$, then by the construction of merge, there is also a transition $x \xrightarrow{a} y$ in $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$. We also need to show that $\lambda_{\mathfrak{M}}(x) \supseteq \lambda_{\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')}(x)$ for all states x in $\mathfrak{M}$. This is immediate from the construction of merge.

□

Proof: To show that $\mathfrak{N} \preceq \mathfrak{M}$ and $\mathfrak{N} \preceq \mathfrak{M}'$ imply $\mathfrak{N} \preceq \mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$, assume there is a simulation $\mathcal{R}$ from $\mathfrak{M}$ to $\mathfrak{N}$ and there is a simulation $\mathcal{R}'$ from $\mathfrak{M}'$ to $\mathfrak{N}$. We need to provide a simulation $\mathcal{R}*$ from $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$ to $\mathfrak{N}$.

Assume the states of $\mathfrak{M}$ and $\mathfrak{M}'$ are disjoint. Define:

$$\mathcal{R}* = \mathcal{R} \cup \mathcal{R}'$$

We need to show that $\mathcal{R}*$ as defined is a simulation from $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$ to $\mathfrak{N}$.

Suppose $x \xrightarrow{a} y$ in $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$ and that $(x, x_2) \in \mathcal{R} \cup \mathcal{R}'$. We need to provide a y_2 such that $x_2 \xrightarrow{a} y_2$ in $\mathfrak{N}$ and $(y, y_2) \in \mathcal{R} \cup \mathcal{R}'$. If $x \xrightarrow{a} y$ in $\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')$, then, from the definition of merge, either $x \xrightarrow{a} y$ in $\mathfrak{M}$ or $x \xrightarrow{a} y$ in $\mathfrak{M}'$. If the former, and given that $\mathcal{R}$ is a simulation from $\mathfrak{M}$ to $\mathfrak{N}$, then there is a y_2 such that $(y, y_2) \in \mathcal{R}$ and $x_2 \xrightarrow{a} y_2$ in $\mathfrak{N}$. But, if $(y, y_2) \in \mathcal{R}$, then also $(y, y_2) \in \mathcal{R} \cup \mathcal{R}'$.

Finally, we need to show that if $(x, y) \in \mathcal{R} \cup \mathcal{R}'$ then

$$\lambda_{\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')}(x) \supseteq \lambda_{\mathfrak{N}}(y)$$

If $(x, y) \in \mathcal{R} \cup \mathcal{R}'$ then either $(x, y) \in \mathcal{R}$ or $(x, y) \in \mathcal{R}'$. Assume the former. Given that $\mathcal{R}$ is a simulation from $\mathfrak{M}$ to $\mathfrak{N}$, we know that if $(x, y) \in \mathcal{R}$, then

$$\lambda_{\mathfrak{M}}(x) \supseteq \lambda_{\mathfrak{N}}(y)$$

Let $\mathfrak{M} = ((S, \rightarrow, \lambda), w)$. If $x \neq w$ (i.e. x is some state other than the start state), then, from the definition of merge, $\lambda_{\mathsf{glb}(\mathfrak{M}, \mathfrak{M}')}(x) = \lambda_{\mathfrak{M}}(x)$. So, given

$\lambda_{\mathfrak{M}} \supseteq \lambda_{\mathfrak{N}}(y)$, $\lambda_{\mathsf{glb}(\mathfrak{M},\mathfrak{M}')}(x) \supseteq \lambda_{\mathfrak{N}}(y)$. If, on the other hand, $x = w$ (i.e. x is the start state of our cathoristic model $\mathfrak{M}$), then, from the definition of merge:

$$\lambda_{\mathsf{glb}(\mathfrak{M},\mathfrak{M}')}(w) = \lambda_{\mathfrak{M}}(w) \cap \lambda_{\mathfrak{M}'}(w')$$

where w' is the start state of $\mathfrak{M}'$. In this case, given $\lambda_{\mathfrak{M}}(w) \supseteq \lambda_{\mathfrak{N}}(y)$ and $\lambda_{\mathfrak{M}'}(w') \supseteq \lambda_{\mathfrak{N}}(y)$, it follows that $\lambda_{\mathfrak{M}}(w) \cap \lambda_{\mathfrak{M}'}(w') \supseteq \lambda_{\mathfrak{N}}(y)$ and hence

$$\lambda_{\mathsf{glb}(\mathfrak{M},\mathfrak{M}')}(w) \supseteq \lambda_{\mathfrak{N}}(y)$$

□

Next, define the least upper bound (lub) of two models as:

$$\begin{aligned}
\mathsf{lub}(\mathfrak{M}, \bot) &= \mathfrak{M} \\
\mathsf{lub}(\bot, \mathfrak{M}) &= \mathfrak{M} \\
\mathsf{lub}((\mathcal{L}, w), (\mathcal{L}', w')) &= \mathsf{lub}_2(\mathcal{L}, \mathcal{L}', (\mathfrak{M}_\top, z), \{(w, w', z)\})
\end{aligned}$$

where $\mathfrak{M}_\top$ is the topmost model ($\mathcal{W} = \{z\}, \rightarrow = \emptyset, \lambda = \{z \mapsto \Sigma\}$) for some state z. lub_2 takes four parameters: the two cathoristic transition systems $\mathcal{L}$ and $\mathcal{L}'$, an accumulator representing the constructed result so far, and a list of state triples (each triple contains one state from each of the two input models plus the state of the accumulated result) to consider next. It is defined as:

$$\begin{aligned}
\mathsf{lub}_2(\mathcal{L}, \mathcal{L}', \mathfrak{M}, \emptyset) &= \mathfrak{M} \\
\mathsf{lub}_2(\mathcal{L}, \mathcal{L}', ((\mathcal{W}, \rightarrow, \lambda), y), \{(w, w', x)\} \cup R) &= \mathsf{lub}_2(\mathcal{L}, \mathcal{L}', ((\mathcal{W} \cup \mathcal{W}', \rightarrow \cup \rightarrow', \lambda'), y), R' \cup R\}
\end{aligned}$$

where:

$$\begin{aligned}
\{(a_i, w_i, w_i') \mid i = 1...n\} &= \mathsf{sharedT}((\mathcal{L}, w), (\mathcal{L}', w')) \\
\mathcal{W}' &= \{x_i \mid i = 1...n\} \\
\rightarrow' &= \{(x, a_i, x_i) \mid i = 1...n\} \\
\lambda' &= \lambda[x \mapsto \lambda(w) \cup \lambda(w)'] \\
R' &= \{(w_i, w_i', x_i) \mid i = 1...n\}
\end{aligned}$$

Here $\lambda[x \mapsto S]$ is the state labelling function that is exactly like λ, except that it maps x to S. Moreover, $\mathsf{sharedT}$ returns the shared transitions between two models, and is defined as:

$$\mathsf{sharedT}(((\mathcal{W}, \rightarrow, \lambda), w)((\mathcal{W}', \rightarrow', \lambda'), w')) = \{(a, x, x') \mid w \xrightarrow{a} x \wedge w' \xrightarrow{a}{}' x'\}$$

If $((S*, \rightarrow *, \lambda*), w*) = ((S, \rightarrow, \lambda), w) \sqcup ((S', \rightarrow', \lambda'), w')$ then define the set $\mathsf{triples}_{\mathsf{lub}}$ as the set of triples $(x, x', x*) \mid x \in S, x' \in S', x* \in S*$ that were used during the construction of lub above. So $\mathsf{triples}_{\mathsf{lub}}$ stores the associations between states in $\mathfrak{M}$, $\mathfrak{M}'$ and $\mathfrak{M} \sqcup \mathfrak{M}'$.

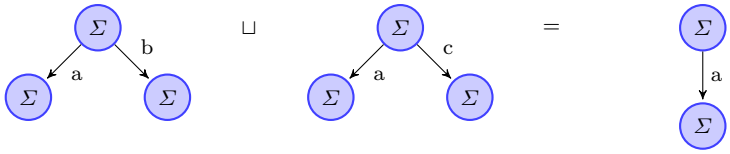

Fig. 9: Example of ⊔

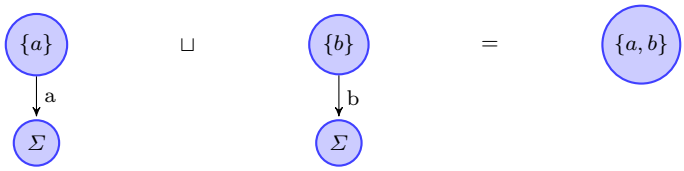

Fig. 10: Example of ⊔

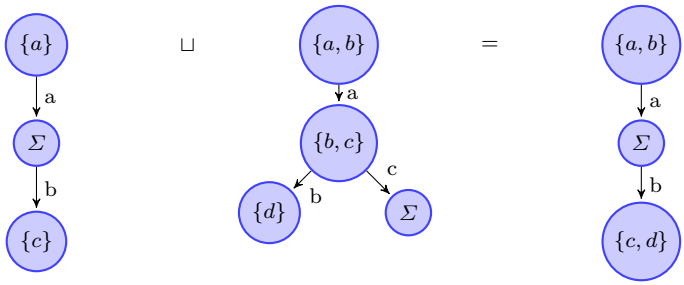

Fig. 11: Example of ⊔

Lemma 4 lub *as defined is the least upper bound*

We will show that:

- $\mathfrak{M} \preceq \mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$ and $\mathfrak{M}' \preceq \mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$
- If $\mathfrak{M} \preceq \mathfrak{N}$ and $\mathfrak{M}' \preceq \mathfrak{N}$, then $\mathsf{lub}(\mathfrak{M}, \mathfrak{M}') \preceq \mathfrak{N}$

If $\mathfrak{M}$ or $\mathfrak{M}'$ are equal to $\bot$, then we just apply the rule that $\bot \preceq m$ for all models m. So let us assume that neither $\mathfrak{M}$ not $\mathfrak{M}'$ are $\bot$.

Proof: To see that $\mathfrak{M} \preceq \mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$, observe that, by construction of lub above, every transition in $\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$ has a matching transition in $\mathfrak{M}$, and every state label in $\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$ is a superset of the corresponding state label in $\mathfrak{M}$.

To show that $\mathfrak{M} \preceq \mathfrak{N}$ and $\mathfrak{M}' \preceq \mathfrak{N}$ together imply $\mathsf{lub}(\mathfrak{M}, \mathfrak{M}') \preceq \mathfrak{N}$, assume a simulation $\mathcal{R}$ from $\mathfrak{N}$ to $\mathfrak{M}$ and a simulation $\mathcal{R}'$ from $\mathfrak{N}$ to $\mathfrak{M}'$. We need to produce a simulation relation $\mathcal{R}*$ from $\mathfrak{N}$ to $\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$. Define

$$\mathcal{R}* = \{(x, y*) \mid \exists y_1.\exists y_2.(x, y_1) \in \mathcal{R}, (x, y_2) \in \mathcal{R}', (y_1, y_2, y*) \in \mathsf{triples}_{\mathsf{lub}}\}$$

In other words, $\mathsf{R}*$ contains the pairs corresponding to the pairs in both R and R'. We just need to show that $\mathsf{R}*$ as defined is a simulation from $\mathfrak{N}$ to $\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$. Assume $(x, x*) \in \mathsf{R}*$ and $x \xrightarrow{a} y$ in $\mathfrak{N}$. We need to produce a $y*$ such that $(x*, y*) \in \mathsf{R}*$ and $x* \xrightarrow{a} y*$ in $\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')$. Given that $\mathcal{R}$ is a simulation from $\mathfrak{N}$ to $\mathfrak{M}$, and that $\mathcal{R}'$ is a simulation from $\mathfrak{N}$ to $\mathfrak{M}'$, we know that there is a pair of states x_1, y_1 in $\mathfrak{M}$ and a pair of states x_2, y_2 in $\mathfrak{M}'$ such that $(x, x_1) \in \mathsf{R}$ and $(x, x_2) \in \mathsf{R}'$ and $x_1 \xrightarrow{a} y_1$ in $\mathfrak{M}$ and $x_2 \xrightarrow{a} y_2$ in $\mathfrak{M}'$. Now, from the construction of lub above, there is a triple $(y_1, y_2, y*) \in \mathsf{triples}_{\mathsf{lub}}$. Now, from the construction of $\mathsf{R}*$ above, $(x*, y*) \in \mathsf{R}*$.

Finally, we need to show that for all states x and y, if $(x, y) \in \mathsf{R}*, \lambda_{\mathfrak{N}}(x) \supseteq \lambda_{\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')}(y)$. Given that $\mathcal{R}$ is a simulation from $\mathfrak{N}$ to $\mathfrak{M}$, and that $\mathcal{R}'$ is a simulation from $\mathfrak{N}$ to $\mathfrak{M}'$, we know that if $(x, y_1) \in \mathsf{R}$, then $\lambda_{\mathfrak{N}}(x) \supseteq \lambda_{\mathfrak{M}}(y_1)$. Similarly, if $(x, y_2) \in \mathsf{R}$, then $\lambda_{\mathfrak{N}}(x) \supseteq \lambda'_{\mathfrak{M}}(y_2)$. Now, from the construction of lub, $\lambda_{\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')}(y*) = \lambda_{\mathfrak{M}}(y_1) \cup \lambda_{\mathfrak{M}}(y_2)$ for all triples $(y_1, y_2, y*) \in \mathsf{triples}_{\mathsf{lub}}$. So $\lambda_{\mathfrak{N}}(x) \supseteq \lambda_{\mathsf{lub}(\mathfrak{M}, \mathfrak{M}')}(y)$, as required. □

6.5 A decision procedure for cathoristic logic

We use the semantic constructions above to provide a quadratic-time decision procedure. The complexity of the decision procedure is an indication that cathoristic logic is useful as a query language in knowledge representation.

Cathoristic logic's lack of connectives for negation, disjunction or implication is the key reason for the efficiency of the decision procedure. Although any satisfiable formula has an infinite number of models, we have shown that the satisfying models form a bounded lattice with a least upper bound. The $\mathsf{simpl}()$ function defined above gives us the least upper bound of all models satisfying an expression. Using this least upper bound, we can calculate entailment by checking a *single model*. To decide whether $\phi \models \psi$, we use the following algorithm.

1. Compute $\mathsf{simpl}(\phi)$.
2. Check if $\mathsf{simpl}(\phi) \models \psi$.

The correctness of this algorithm is given by the follow theorem.

Theorem 3 *The following are equivalent:*

1. *For all cathoristic models* $\mathfrak{M}$, $\mathfrak{M} \models \phi$ *implies* $\mathfrak{M} \models \psi$.
2. $\mathsf{simpl}(\phi) \models \psi$.

Proof: The implication from (1) to (2) is trivial because $\mathsf{simpl}(\phi) \models \phi$ by construction.

For the reverse direction, we make use of the following lemma (proved in the Appendix):

Lemma 5 *If* $\mathfrak{M} \models \phi$ *then* $\mathfrak{M} \preceq \mathsf{simpl}(\phi)$.

With Lemma 5 in hand, the proof of Theorem 3 is straightforward. Assume $\mathfrak{M} \models \phi$. We need to show $\mathfrak{M} \models \psi$. Now if $\mathfrak{M} \models \phi$ then $\mathfrak{M} \preceq \mathsf{simpl}(\phi)$ (by Lemma 5). Further, if $\mathfrak{M}' \models \xi$ and $\mathfrak{M} \preceq \mathfrak{M}'$ then $\mathfrak{M} \models \xi$ by Theorem 1. So, substituting ψ for ξ and $\mathsf{simpl}(\phi)$ for $\mathfrak{M}'$, it follows that $\mathfrak{M} \models \psi$. □

Construction of $\mathsf{simpl}(\phi)$ is quadratic in the size of ϕ, and computing whether a model satisfies ψ is of order $|\psi| \times |\phi|$, so computing whether $\phi \models \psi$ is quadratic time.

6.6 Incompatibility semantics

One of cathoristic logic's unusual features is that it satisfies Brandom's incompatibility semantics constraint, *even though it has no negation operator.* In this section, we formalise what this means, and prove it.

Define the *incompatibility set of* ϕ as:

$$\mathcal{I}(\phi) = \{\psi \mid \forall \mathfrak{M}.\mathfrak{M} \not\models \phi \wedge \psi\}$$

The reason why Brandom introduces the incompatibility set [14] is that he wants to use it define *semantic content*:

> Here is a semantic suggestion: represent the propositional content expressed by a sentence with the set of sentences that express propositions incompatible with it[15].

Now if the propositional content of a claim determines its logical consequences, and the propositional content is identified with the incompatibility set, then the incompatibility set must determine the logical consequences. A logic satisfies Brandom's *incompatibility semantics constraint* if

$$\phi \models \psi \quad \text{iff} \quad \mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$$

Not all logics satisfy this property. Brandom has shown that first-order logic and the modal logic S5 satisfy the incompatibility semantics property. Hennessy-Milner logic satisfies it, but Hennessy-Milner logic without negation does not. Cathoristic logic is the simplest logic we have found that satisfies the property. To establish the incompatibility semantics constraint for cathoristic logic, we

[14] Brandom [8] defines incompatibility slightly differently: he defines the set of *sets* of formulae which are incompatible with a *set* of formulae. But in cathoristic logic, if a set of formulae is incompatible, then there is an incompatible subset of that set with exactly two members. So we can work with the simpler definition in the text above.

[15] [8] p.123.

need to define a related incompatibility function on models. $\mathcal{J}(\mathfrak{M})$ is the set of models that are incompatible with $\mathfrak{M}$:

$$\mathcal{J}(\mathfrak{M}) = \{\mathfrak{M}_2 \mid \mathfrak{M} \sqcap \mathfrak{M}_2 = \bot\}$$

We shall make use of two lemmas, proved in Appendix B:

Lemma 6 *If* $\phi \models \psi$ *then* $\mathsf{simpl}(\phi) \preceq \mathsf{simpl}(\psi)$

Lemma 7 $\mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$ *implies* $\mathcal{J}(\mathsf{simpl}(\psi)) \subseteq \mathcal{J}(\mathsf{simpl}(\phi))$

Theorem 4 $\phi \models \psi$ *iff* $\mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$

Proof: Left to right: Assume $\phi \models \psi$ and $\xi \in \mathcal{I}(\psi)$. We need to show $\xi \in \mathcal{I}(\phi)$. By the definition of $\mathcal{I}$, if $\xi \in \mathcal{I}(\psi)$ then $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\psi) = \bot$. If $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\psi) = \bot$, then either

- $\mathsf{simpl}(\xi) = \bot$
- $\mathsf{simpl}(\psi) = \bot$
- Neither $\mathsf{simpl}(\xi)$ nor $\mathsf{simpl}(\psi)$ are $\bot$, but $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\psi) = \bot$.

If $\mathsf{simpl}(\xi) = \bot$, then $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\phi) = \bot$ and we are done. If $\mathsf{simpl}(\psi) = \bot$, then as $\phi \models \psi$, by Lemma 6, $\mathsf{simpl}(\phi) \preceq \mathsf{simpl}(\psi)$. Now the only model that is $\preceq \bot$ is $\bot$ itself, so $\mathsf{simpl}(\phi) = \bot$. Hence $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\phi) = \bot$, and we are done. The interesting case is when neither $\mathsf{simpl}(\xi)$ nor $\mathsf{simpl}(\psi)$ are $\bot$, but $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\psi) = \bot$. Then (by the definition of $\mathsf{consistent}$ in Section 6.4), either $\mathsf{out}(\mathsf{simpl}(\xi)) \not\subseteq \lambda(\mathsf{simpl}(\psi))$ or $\mathsf{out}(\mathsf{simpl}(\psi)) \not\subseteq \lambda(\mathsf{simpl}(\xi))$. In the first sub-case, if $\mathsf{out}(\mathsf{simpl}(\xi)) \not\subseteq \lambda(\mathsf{simpl}(\psi))$, then there is some action a such that $\xi \models \langle a \rangle \top$ and $a \notin \lambda(\mathsf{simpl}(\psi))$. If $a \notin \lambda(\mathsf{simpl}(\psi))$ then $\psi \models !A$ where $a \notin A$. Now $\phi \models \psi$, so $\phi \models !A$. In other words, ϕ also entails the A-restriction that rules out the a transition. So $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\phi) = \bot$ and $\xi \in \mathcal{I}(\phi)$. In the second sub-case, $\mathsf{out}(\mathsf{simpl}(\psi)) \not\subseteq \lambda(\mathsf{simpl}(\xi))$. Then there is some action a such that $\psi \models \langle a \rangle \top$ and $a \notin \lambda(\mathsf{simpl}(\xi))$. If $a \notin \lambda(\mathsf{simpl}(\xi))$ then $\xi \models !A$ where $a \notin A$. But if $\psi \models \langle a \rangle \top$ and $\phi \models \psi$, then $\phi \models \langle a \rangle \top$ and ϕ is also incompatible with ξ's A-restriction. So $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\phi) = \bot$ and $\xi \in \mathcal{I}(\phi)$.

Right to left: assume, for reductio, that $\mathfrak{M} \models \phi$ and $\mathfrak{M} \not\models \psi$. we will show that $\mathcal{I}(\psi) \not\subseteq \mathcal{I}(\phi)$. Assume $\mathfrak{M} \models \phi$ and $\mathfrak{M} \not\models \psi$. We will construct another model $\mathfrak{M}_2$ such that $\mathfrak{M}_2 \in \mathcal{J}(\mathsf{simpl}(\psi))$ but $\mathfrak{M}_2 \notin \mathcal{J}(\mathsf{simpl}(\phi))$. This will entail, via Lemma 7, that $\mathcal{I}(\psi) \not\subseteq \mathcal{I}(\phi)$.

If $\mathfrak{M} \not\models \psi$, then there is a formula ψ' that does not contain $\wedge$ such that $\psi \models \psi'$ and $\mathfrak{M} \not\models \psi'$. ψ' must be either of the form (i) $\langle a_1 \rangle ... \langle a_n \rangle \top$ (for $n > 0$) or (ii) of the form $\langle a_1 \rangle ... \langle a_n \rangle\, !\{A\}$ where $A \subseteq \mathcal{S}$ and $n >= 0$.

In case (i), there must be an i between 0 and n such that $\mathfrak{M} \models \langle a_1 \rangle ... \langle a_i \rangle \top$ but $\mathfrak{M} \not\models \langle a_1 \rangle ... \langle a_{i+1} \rangle \top$. We need to construct another model $\mathfrak{M}_2$ such that $\mathfrak{M}_2 \sqcap \mathsf{simpl}(\psi) = \bot$, but $\mathfrak{M}_2 \sqcap \mathsf{simpl}(\phi) \neq \bot$. Letting $\mathfrak{M} = ((\mathcal{W}, \rightarrow, \lambda), w)$, then $\mathfrak{M} \models \langle a_1 \rangle ... \langle a_i \rangle \top$ implies that there is at least one sequence of states of the form $w, w_1, ..., w_i$ such that $w \xrightarrow{a_1} w_1 \rightarrow ... \xrightarrow{a_i} w_i$. Now let $\mathfrak{M}_2$ be just like $\mathfrak{M}$ but with additional transition-restrictions on each w_i that it not include a_{i+1}.

In other words, $\lambda_{\mathfrak{M}_2}(w_i) = \lambda_{\mathfrak{M}}(w_i) - \{a_{i+1}\}$ for all w_i in sequences of the form $w \xrightarrow{a_1} w_1 \to ... \xrightarrow{a_i} w_i$. Now $\mathfrak{M}_2 \sqcap \mathsf{simpl}(\psi) = \bot$ because of the additional transition restriction we added to $\mathfrak{M}_2$, which rules out $\langle a_1\rangle...\langle a_{i+1}\rangle\top$, and a-fortiori ψ. But $\mathfrak{M}_2 \sqcap \mathsf{simpl}(\phi) \neq \bot$, because $\mathfrak{M} \models \phi$ and $\mathfrak{M}_2 \preceq \mathfrak{M}$ together imply $\mathfrak{M}_2 \models \phi$. So $\mathfrak{M}_2$ is indeed the model we were looking for, that is incompatible with $\mathsf{simpl}(\psi)$ while being compatible with $\mathsf{simpl}(\phi)$.

In case (ii), $\mathfrak{M} \models \langle a_1\rangle...\langle a_n\rangle\top$ but $\mathfrak{M} \not\models \langle a_1\rangle...\langle a_n\rangle!A$ for some $A \subset \mathcal{S}$. We need to produce a model $\mathfrak{M}_2$ that is incompatible with $\mathsf{simpl}(\psi)$ but not with $\mathsf{simpl}(\phi)$. Given that $\mathfrak{M} \models \langle a_1\rangle...\langle a_n\rangle\top$, there is a sequence of states $w, w_1, ..., w_n$ such that $w \xrightarrow{a_1} w_1 \to ... \xrightarrow{a_i} w_n$. Let $\mathfrak{M}_2$ be the model just like $\mathfrak{M}$ except it has an additional transition from each such w_n with an action $a \notin A$. Clearly, $\mathfrak{M}_2 \sqcap \mathsf{simpl}(\psi') = \bot$ because of the additional a-transition, and given that $\psi \models \psi'$, it follows that $\mathfrak{M}_2 \sqcap \mathsf{simpl}(\psi) = \bot$. Also, $\mathfrak{M}_2 \sqcap \mathsf{simpl}(\phi) \neq \bot$, because $\mathfrak{M}_2 \preceq \mathfrak{M}$ and $\mathfrak{M} \models \phi$.

□

7 Inference Rules

$$\frac{-}{\phi \vdash \phi}\ \textsc{Id} \qquad \frac{-}{\phi \vdash \top}\ \top\text{-}\textsc{Right} \qquad \frac{-}{\bot \vdash \phi}\ \bot\text{-}\textsc{Left} \qquad \frac{\phi \vdash \psi \quad \psi \vdash \xi}{\phi \vdash \xi}\ \textsc{Trans}$$

$$\frac{\phi \vdash \psi}{\phi \wedge \xi \vdash \psi}\ \wedge\text{-}\textsc{Left 1} \qquad \frac{\phi \vdash \psi}{\xi \wedge \phi \vdash \psi}\ \wedge\text{-}\textsc{Left 2} \qquad \frac{\phi \vdash \psi \quad \phi \vdash \xi}{\phi \vdash \psi \wedge \xi}\ \wedge\text{-}\textsc{Right}$$

$$\frac{a \notin A}{!A \wedge \langle a \rangle \phi \vdash \bot}\ \bot\text{-}\textsc{Right 1} \qquad \frac{-}{\langle a \rangle \bot \vdash \bot}\ \bot\text{-}\textsc{Right 2} \qquad \frac{\phi \vdash !A \quad A \subset A'}{\phi \vdash !A'}\ !\text{-}\textsc{Right 1}$$

$$\frac{\phi \vdash !A \quad \phi \vdash !B}{\phi \vdash !(A \cap B)}\ !\text{-}\textsc{Right 2} \qquad \frac{\phi \vdash \psi}{\langle a \rangle \phi \vdash \langle a \rangle \psi}\ \textsc{Normal} \qquad \frac{\phi \vdash \langle a \rangle \psi \wedge \langle a \rangle \xi}{\phi \vdash \langle a \rangle (\psi \wedge \xi)}\ \textsc{Det}$$

Fig. 12: Proof rules.

We now present the inference rules for cathoristic logic. There are no axioms.

Definition 12 Judgements are of the following form.

$$\phi \vdash \psi.$$

We also write $\vdash \phi$ as a shorthand for $\top \vdash \phi$. Figure 12 presents all proof rules.

Note that ϕ and ψ are single formulae, not sequents. By using single formulae, we can avoid structural inference rules. The proof rules can be grouped in two parts: standard rules and rules unique to cathoristic logic. Standard rules are [Id], [$\top$-Right], [$\bot$-Left], [Trans], [$\wedge$-Left 1], [$\wedge$-Left 2] and [$\wedge$-Right]. They hardly need explanation as they are variants of familiar rules for propositional logic, see e.g. [30,32]. We now explain the rules that give cathoristic logic its distinctive properties.

The rule [$\bot$-Right 1] captures the core exclusion property of the tantum !: for example if $A = \{male, female\}$ then $\langle orange \rangle \phi$ is incompatible with $!A$. Thus $!A \wedge \langle orange \rangle \phi$ must be false.

The rule [$\bot$-Right 2] expresses that falsity is 'global' and cannot be suppressed by the modalities. For example $\langle orange \rangle \bot$ is false, simply because $\bot$ is already false.

[Normal] enables us to prefix an inference with a may-modality. This rule can also be stated in the the following more general form:

$$\frac{\phi_1 \wedge ... \wedge \phi_n \vdash \psi}{\langle a \rangle \phi_1 \wedge ... \wedge \langle a \rangle \phi_n \vdash \langle a \rangle \psi}\ \textsc{Normal-Multi}$$

But it is not necessary because [Normal-Multi] is derivable from [Normal] as we show in the examples below.

7.1 Example inferences

We prove that we can use $\phi \wedge \psi \vdash \xi$ to derive $\langle a\rangle\phi \wedge \langle a\rangle\psi \vdash \langle a\rangle\xi$:

$$\dfrac{\dfrac{\langle a\rangle\phi \wedge \langle a\rangle\psi \vdash \langle a\rangle\phi \wedge \langle a\rangle\psi}{\langle a\rangle\phi \wedge \langle a\rangle\psi \vdash \langle a\rangle(\phi \wedge \psi)}\,\textsc{Det} \qquad \dfrac{\phi \wedge \psi \vdash \xi}{\langle a\rangle(\phi \wedge \psi) \vdash \langle a\rangle\xi}\,\textsc{Normal}}{\langle a\rangle\phi \wedge \langle a\rangle\psi \vdash \langle a\rangle\xi}\,\textsc{Trans}$$

Figure 13 demonstrates how to infer $\langle a\rangle!\{b,c\} \wedge \langle a\rangle!\{c,d\} \vdash \langle a\rangle!\{c\}$ and $\langle a\rangle!\{b\} \wedge \langle a\rangle\langle c\rangle\top \vdash \langle d\rangle\top$.

7.2 !-Left and !-Right

The rules [!-Right 1, !-Right 2] jointly express how the subset relation $\subseteq$ on sets of actions relates to provability. Why don't we need a corresponding rule !-Left for strengthening ! on the left hand side?

$$\dfrac{\phi \wedge\, !A \vdash \psi \quad A' \subseteq A}{\phi \wedge\, !A' \vdash \psi}\,\text{!-}\textsc{Left}$$

The reason is that [!-Left] can be derived as follows.

$$\dfrac{\dfrac{\phi \wedge\, !A' \vdash \phi \wedge\, !A' \qquad A' \subseteq A}{\phi \wedge\, !A' \vdash \phi \wedge\, !A}\,\text{!-}\textsc{Right 1} \qquad \phi \wedge\, !A \vdash \psi}{\phi \wedge\, !A' \vdash \psi}\,\textsc{Trans}$$

Readers familiar with object-oriented programming will recognise [!-Left] as contravariant and [!-Right 1] as covariant subtyping. Honda [18] develops a full theory of subtyping based on similar ideas. All three rules embody the intuition that whenever $A \subseteq A'$ then asserting that $!A'$ is as strong as, or a stronger statement than $!A$. [!-Left] simply states that we can always strengthen our premise, while [!-right 1] allows us to weaken the conclusion.

7.3 Characteristic formulae

In order to prove completeness, below, we need the notion of a *characteristic formula* of a model. The function $\mathsf{simpl}(\cdot)$ takes a formula as argument and returns the least upper bound of the satisfying models. Characteristic formulae go the other way: given a model $\mathfrak{M}$, $\mathsf{char}(\mathfrak{M})$ is the logically weakest formula that describes that model.

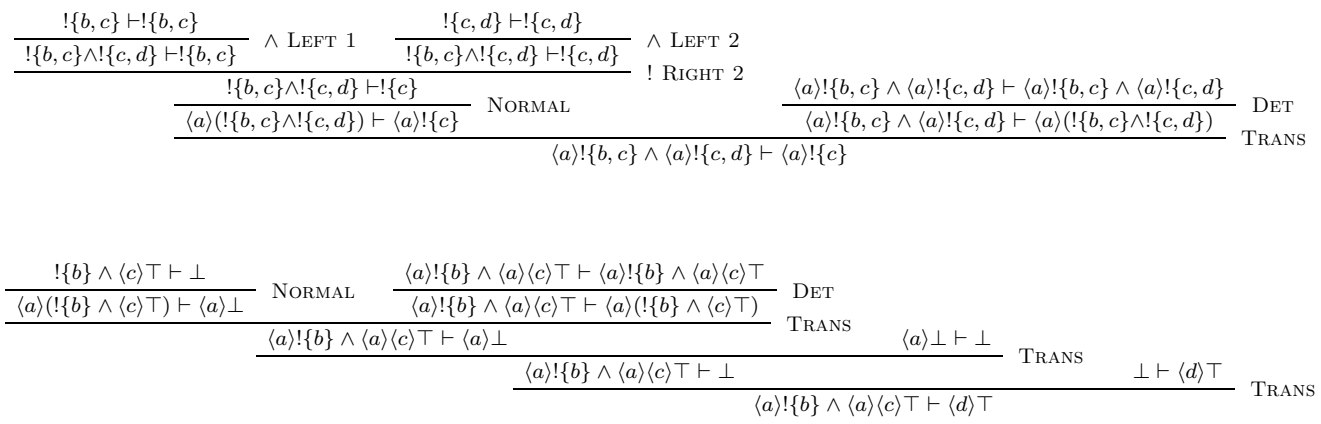
!{b, c} ⊢ !{b, c}
!{b, c}∧!{c, d} ⊢ !{b, c}
∧ Left 1
!{c, d} ⊢ !{c, d}
!{b, c}∧!{c, d} ⊢ !{c, d}
∧ Left 2
! Right 2
!{b, c}∧!{c, d} ⊢ !{c}
⟨a⟩(!{b, c}∧!{c, d}) ⊢ ⟨a⟩!{c}
Normal
⟨a⟩!{b, c} ∧ ⟨a⟩!{c, d} ⊢ ⟨a⟩!{b, c} ∧ ⟨a⟩!{c, d}
⟨a⟩!{b, c} ∧ ⟨a⟩!{c, d} ⊢ ⟨a⟩(!{b, c}∧!{c, d})
Det
⟨a⟩!{b, c} ∧ ⟨a⟩!{c, d} ⊢ ⟨a⟩!{c}
Trans
!{b} ∧ ⟨c⟩⊤ ⊢ ⊥
⟨a⟩(!{b} ∧ ⟨c⟩⊤) ⊢ ⟨a⟩⊥
Normal
⟨a⟩!{b} ∧ ⟨a⟩⟨c⟩⊤ ⊢ ⟨a⟩!{b} ∧ ⟨a⟩⟨c⟩⊤
⟨a⟩!{b} ∧ ⟨a⟩⟨c⟩⊤ ⊢ ⟨a⟩(!{b} ∧ ⟨c⟩⊤)
Det
⟨a⟩!{b} ∧ ⟨a⟩⟨c⟩⊤ ⊢ ⟨a⟩⊥
Trans
⟨a⟩⊥ ⊢ ⊥
⟨a⟩!{b} ∧ ⟨a⟩⟨c⟩⊤ ⊢ ⊥
Trans
⊥ ⊢ ⟨d⟩⊤
⟨a⟩!{b} ∧ ⟨a⟩⟨c⟩⊤ ⊢ ⟨d⟩⊤
Trans

Fig. 13: Derivations of $\langle a\rangle!\{b,c\} \wedge \langle a\rangle!\{c,d\} \vdash \langle a\rangle!\{c\}$ (top) and $\langle a\rangle!\{b\} \wedge \langle a\rangle\langle c\rangle\top \vdash \langle d\rangle\top$ (bottom).

Definition 13 Let $\mathfrak{M}$ be a cathoristic model that is a tree.

$$\begin{aligned} \mathsf{char}(\bot) &= \langle a\rangle\top \wedge !\emptyset \text{ for some fixed action } a \in \Sigma \\ \mathsf{char}(\mathfrak{M}, w) &= \mathsf{bang}(\mathfrak{M}, w) \wedge \bigwedge_{w \xrightarrow{a} w'} \langle a\rangle \mathsf{char}(\mathfrak{M}, w') \end{aligned}$$

Note that $\bot$ requires a particular action $a \in \Sigma$. This is why we required, in Section 3.1, that Σ is non-empty.

The functions $\mathsf{bang}(\cdot)$ on models are given by the following clauses.

$$\mathsf{bang}((S, \to, \lambda), w) = \begin{cases} \top & \text{if } \lambda(w) = \Sigma \\ !\,\lambda(w) & \text{otherwise} \end{cases}$$

Note that $\mathsf{char}(\mathfrak{M})$ is finite if $\mathfrak{M}$ contains no cycles and if $\lambda(x)$ is either Σ or finite for all states x. We state without proof that $\mathsf{simpl}(\cdot)$ and $\mathsf{char}(\cdot)$ are inverses of each other (for tree models $\mathfrak{M}$) in that:

- $\mathsf{simpl}(\mathsf{char}(\mathfrak{M})) \simeq \mathfrak{M}$.
- $\models \mathsf{char}(\mathsf{simpl}(\phi))$ iff $\models \phi$.

7.4 Soundness and completeness

Theorem 5 *The rules in Figure 12 are sound and complete:*

1. *(Soundness)* $\phi \vdash \psi$ *implies* $\phi \models \psi$.
2. *(Completeness)* $\phi \models \psi$ *implies* $\phi \vdash \psi$.

Soundness is immediate from the definitions. To prove completeness we will show that $\phi \models \psi$ implies there is a derivation of $\phi \vdash \psi$. Our proof will make use of two key facts (proved in Sections 7.5.1 and 7.5.2 below):

Lemma 8 *If* $\mathfrak{M} \models \phi$ *then* $\mathsf{char}(\mathfrak{M}) \vdash \phi$.

Lemma 9 *For all formulae* ϕ*, we can derive* $\phi \vdash \mathsf{char}(\mathsf{simpl}(\phi))$.

Lemma 8 states that, if ϕ is satisfied by a model, then there is a proof that the characteristic formula describing that model entails ϕ. In Lemma 9, $\mathsf{simpl}(\phi)$ is the simplest model satisfying ϕ, and $\mathsf{char}(\mathfrak{M})$ is the simplest formula describing m, so $\mathsf{char}(\mathsf{simpl}(\phi))$ is a simplified form of ϕ. This lemma states that cathoristic logic has the inferential capacity to transform any proposition into its simplified form.

With these two lemmas in hand, the proof of completeness is straightforward. Assume $\phi \models \psi$. Then all models which satisfy ϕ also satisfy ψ. In particular, $\mathsf{simpl}(\phi) \models \psi$. Then $\mathsf{char}(\mathsf{simpl}(\phi)) \vdash \psi$ by Lemma 8. But we also have, by Lemma 9, $\phi \vdash \mathsf{char}(\mathsf{simpl}(\phi))$. So by transitivity, we have $\phi \vdash \psi$.

7.5 Proofs of Lemmas 8, 9 and 10

7.5.1 Proof of Lemma 8

If $\mathfrak{M} \models \phi$ then $\mathsf{char}(\mathfrak{M}) \vdash \phi$.
We proceed by induction on ϕ.

Case ϕ is $\top$. Then we can prove $\mathsf{char}(\mathfrak{M}) \vdash \phi$ immediately using axiom [$\top$ RIGHT.

Case ϕ is $\psi \wedge \psi'$. By the induction hypothesis, $\mathsf{char}(\mathfrak{M}) \vdash \psi$ and $\mathsf{char}(\mathfrak{M}) \vdash \psi'$. The proof of $\mathsf{char}(\mathfrak{M}) \vdash \psi \wedge \psi'$ follows immediately using [$\wedge$ RIGHT.

Case ϕ is $\langle a \rangle \psi$. If $\mathfrak{M} \models \langle a \rangle \psi$, then either $\mathfrak{M} = \bot$ or $\mathfrak{M}$ is a model of the form $(\mathcal{L}, w)$.

Subcase $\mathfrak{M} = \bot$. In this case, $\mathsf{char}(\mathfrak{M}) = \mathsf{char}(\bot) = \bot$. (Recall, that we are overloading $\bot$ to mean both the model at the bottom of our lattice and a formula (such as $\langle a \rangle \top \wedge !\emptyset$) which is always false). In this case, $\mathsf{char}(\bot) \vdash \langle a \rangle \psi$ using [$\bot$ LEFT.

Subcase m is a model of the form $(\mathcal{L}, w)$. Given $\mathfrak{M} \models \langle a \rangle \psi$, and that $\mathfrak{M}$ is a model of the form $(\mathcal{L}, w)$, we know that:

$$(\mathcal{L}, w) \models \langle a \rangle \psi$$

From the satisfaction clause for $\langle a \rangle$, it follows that:

$$\exists w' \text{ such that } w \xrightarrow{a} w' \text{ and } (\mathcal{L}, w') \models \psi$$

By the induction hypothesis:

$$\mathsf{char}((\mathcal{L}, w')) \vdash \psi$$

Now by [NORMAL]:

$$\langle a \rangle \mathsf{char}((\mathcal{L}, w')) \vdash \langle a \rangle \psi$$

Using repeated application of [$\wedge$ LEFT], we can show:

$$\mathsf{char}((\mathcal{L}, w)) \vdash \langle a \rangle \mathsf{char}((\mathcal{L}, w'))$$

Finally, using [TRANS], we derive:

$$\mathsf{char}((\mathcal{L}, w)) \vdash \langle a \rangle \psi$$

Case ϕ is $!\psi$. If $(\mathcal{L}, w) \models !A$, then $\lambda(w) \subseteq A$. Then $\mathsf{char}((\mathcal{L}, w)) =! \lambda(w) \wedge \phi$. Now we can prove $! \lambda(w) \wedge \phi \vdash !A$ using [! RIGHT 1] and repeated applications of [$\wedge$ LEFT].

7.5.2 Proof of Lemma 9

Now we prove Lemma 9: for all formulae ϕ, we can derive $\phi \vdash \mathsf{char}(\mathsf{simpl}(\phi))$.

Proof: Induction on ϕ.

Case ϕ is $\top$. Then we can prove $\top \vdash \top$ using either [$\top$ RIGHT] or [ID].

Case ϕ is $\psi \wedge \psi'$. By the induction hypothesis, $\psi \vdash \mathsf{char}(\mathsf{simpl}(\psi))$ and $\psi' \vdash \mathsf{char}(\mathsf{simpl}(\psi'))$. Using [$\wedge$ LEFT] and [$\wedge$ RIGHT], we can show:

$$\psi \wedge \psi' \vdash \mathsf{char}(\mathsf{simpl}(\psi)) \wedge \mathsf{char}(\mathsf{simpl}(\psi'))$$

In order to continue the proof, we need the following lemma, proven in the next subsection.

Lemma 10 *For all cathoristic models $\mathfrak{M}$ and $\mathfrak{M}_2$ that are trees,* $\mathsf{char}(\mathfrak{M}) \wedge \mathsf{char}(\mathfrak{M}_2) \vdash \mathsf{char}(\mathfrak{M} \sqcap \mathfrak{M}_2)$.

From Lemma 10 (substituting $\mathsf{simpl}(\psi)$ for $\mathfrak{M}$ and $\mathsf{simpl}(\psi')$ for $\mathfrak{M}_2$, and noting that $\mathsf{simpl}()$ always produces acyclic models), it follows that:

$$\mathsf{char}(\mathsf{simpl}(\psi)) \wedge \mathsf{char}(\mathsf{simpl}(\psi')) \vdash \mathsf{char}(\mathsf{simpl}(\psi \wedge \psi'))$$

Our desired result follows using [TRANS].

Case ϕ is $\langle a \rangle \psi$. By the induction hypothesis, $\psi \vdash \mathsf{char}(\mathsf{simpl}(\psi))$. Now there are two sub-cases to consider, depending on whether or not $\mathsf{char}(\mathsf{simpl}(\psi)) = \bot$.

Subcase $\mathsf{char}(\mathsf{simpl}(\psi)) = \bot$. In this case, $\mathsf{char}(\mathsf{simpl}(\langle a \rangle \psi))$ also equals $\bot$. By the induction hypothesis:

$$\psi \vdash \bot$$

By [NORMAL]:

$$\langle a \rangle \psi \vdash \langle a \rangle \bot$$

By [$\bot$ RIGHT 2]:

$$\langle a \rangle \bot \vdash \bot$$

The desired proof that:

$$\langle a \rangle \psi \vdash \bot$$

follows by [TRANS].

Subcase $\mathsf{char}(\mathsf{simpl}(\psi)) \neq \bot$. By the induction hypothesis, $\psi \vdash \mathsf{char}(\mathsf{simpl}(\psi))$. So, by [NORMAL]:

$$\langle a \rangle \psi \vdash \langle a \rangle \mathsf{char}(\mathsf{simpl}(\psi))$$

The desired conclusion follows from noting that:

$$\langle a \rangle \mathsf{char}(\mathsf{simpl}(\psi)) = \mathsf{char}(\mathsf{simpl}(\langle a \rangle \psi))$$

Case ϕ is $!A$. If ϕ is $!A$, then $\mathsf{char}(\mathsf{simpl}(\phi))$ is $!A \wedge \top$. We can prove $!A \vdash !A \wedge \top$ using [$\wedge$ RIGHT], [$\top$ RIGHT] and [ID]. □

7.5.3 Proof of Lemma 10

We can now finish the proof of Lemma 9 by giving the missing proof of Lemma 10.

Proof: There are two cases to consider, depending on whether or not $(\mathfrak{M} \sqcap \mathfrak{M}_2) = \bot$.

Case $(\mathfrak{M} \sqcap \mathfrak{M}_2) = \bot$. If $(\mathfrak{M} \sqcap \mathfrak{M}_2) = \bot$, there are three possibilities:

- $\mathfrak{M} = \bot$
- $\mathfrak{M}_2 = \bot$
- Neither $\mathfrak{M}$ nor $\mathfrak{M}_2$ are $\bot$, but together they are incompatible.

If either $\mathfrak{M}$ or $\mathfrak{M}_2$ is $\bot$, then the proof is a simple application of [Id] followed by [$\wedge$ Left].

Next, let us consider the case where neither $\mathfrak{M}$ nor $\mathfrak{M}_2$ are $\bot$, but together they are incompatible. Let $\mathfrak{M} = (\mathcal{L}, w_1)$ and $\mathfrak{M}' = (\mathcal{L}', w_1')$. If $\mathfrak{M} \sqcap \mathfrak{M}_2 = \bot$, then there is a finite sequence of actions $a_1, ..., a_{n-1}$ such that both $\mathfrak{M}$ and $\mathfrak{M}'$ satisfy $\langle a_1 \rangle ... \langle a_{n-1} \rangle \top$, but they disagree about the state-labelling on the final state of this chain. In other words, there is a b-transition from the final state in $\mathfrak{M}$ which is ruled-out by the λ' state-labelling in $\mathfrak{M}'$. So there is a set of states $w_1, ..., w_1', ...$ and a finite set X of actions such that:

- $w_1 \xrightarrow{a_1} w_2 \xrightarrow{a_2} ... \xrightarrow{a_{n-1}} w_n$.
- $w_1' \xrightarrow{a_1} w_2' \xrightarrow{a_2} ... \xrightarrow{a_{n-1}} w_n'$.
- $w_n \xrightarrow{b} w_{n+1}$.
- $\lambda'(w_n') = X$ with $b \notin X$.

Now it is easy to show, using [$\wedge$ Left], that

$$\mathsf{char}(\mathfrak{M}) \vdash \langle a_1 \rangle ... \langle a_{n-1} \rangle \langle b \rangle \top$$
$$\mathsf{char}(\mathfrak{M}') \vdash \langle a_1 \rangle ... \langle a_{n-1} \rangle !X$$

Now using [$\wedge$ Left] and [$\wedge$ Right]:

$$\mathsf{char}(\mathfrak{M}) \wedge \mathsf{char}(\mathfrak{M}') \vdash \langle a_1 \rangle ... \langle a_{n-1} \rangle \langle b \rangle \top \wedge \vdash \langle a_1 \rangle ... \langle a_{n-1} \rangle !X$$

Now using [Det]:

$$\mathsf{char}(\mathfrak{M}) \wedge \mathsf{char}(\mathfrak{M}') \vdash \langle a_1 \rangle ... \langle a_{n-1} \rangle (\langle b \rangle \top \wedge !X)$$

Now, using [$\bot$ Right 1]:

$$\langle b \rangle \top \wedge !X \vdash \bot$$

Using $n-1$ applications of [$\bot$ Right 2]:

$$\langle a_1 \rangle ... \langle a_{n-1} \rangle (\langle b \rangle \top \wedge !X) \vdash \bot$$

Finally, using [Trans], we derive:

$$\mathsf{char}(\mathfrak{M}) \wedge \mathsf{char}(\mathfrak{M}') \vdash \bot$$

Case $(\mathfrak{M} \sqcap \mathfrak{M}_2) \neq \bot$. From the construction of merge, if $\mathfrak{M}$ and $\mathfrak{M}'$ are acyclic, then $\mathfrak{M} \sqcap \mathfrak{M}'$ is also acyclic. If $\mathfrak{M} \sqcap \mathfrak{M}'$ is acyclic, then $\mathsf{char}(\mathfrak{M} \sqcap \mathfrak{M}')$ is equivalent to a set Γ of sentences of one of two forms:

$$\langle a_1 \rangle ... \langle a_n \rangle \top \qquad\qquad \langle a_1 \rangle ... \langle a_n \rangle !X$$

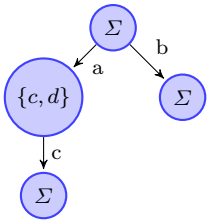

Fig. 14: Example of ⊓

For example, if $\mathfrak{M} \sqcap \mathfrak{M}'$ is as in Figure 14, then

$$\mathsf{char}(\mathfrak{M} \sqcap \mathfrak{M}') = \langle a \rangle (!\{c, d\} \wedge \langle c \rangle \top) \wedge \langle b \rangle \top$$

This is equivalent to the set Γ of sentences:

$$\langle a \rangle \langle c \rangle \top \qquad \langle b \rangle \top \qquad \langle a \rangle !\{c, d\}$$

Now using [$\wedge$ Right] and [Det] we can show that

$$\bigwedge_{\phi \in \Gamma} \phi \vdash \mathsf{char}(\mathfrak{M} \sqcap \mathfrak{M}')$$

We know that for all $\phi \in \Gamma$

$$\mathfrak{M} \sqcap \mathfrak{M}' \models \phi$$

We just need to show that:

$$\mathsf{char}(\mathfrak{M}) \wedge \mathsf{char}(\mathfrak{M}') \vdash \phi$$

Take any $\phi \in \Gamma$ of the form $\langle a_1 \rangle ... \langle a_n \rangle !X$ for some finite $X \subseteq \Sigma$. (The case where ϕ is of the form $\langle a_1 \rangle ... \langle a_n \rangle \top$ is very similar, but simpler). If $\mathfrak{M} \sqcap \mathfrak{M}' \models \langle a_1 \rangle ... \langle a_n \rangle !X$ then either:

1. $\mathfrak{M} \models \langle a_1 \rangle ... \langle a_n \rangle !X$ but $\mathfrak{M}' \not\models \langle a_1 \rangle ... \langle a_n \rangle \top$
2. $\mathfrak{M}' \models \langle a_1 \rangle ... \langle a_n \rangle !X$ but $\mathfrak{M} \not\models \langle a_1 \rangle ... \langle a_n \rangle \top$
3. $\mathfrak{M} \models \langle a_1 \rangle ... \langle a_n \rangle !X_1$ and $\mathfrak{M}' \models \langle a_1 \rangle ... \langle a_n \rangle !X_2$ and $X_1 \cap X_2 \subseteq X$

In the first two cases, showing $\mathsf{char}(\mathfrak{M}) \wedge \mathsf{char}(\mathfrak{M}') \vdash \phi$ is just a matter of repeated application of [$\wedge$ Left] and [$\wedge$ Right]. In the third case, let $\mathfrak{M} = (\mathcal{L}, w_1)$ and $\mathfrak{M}' = (\mathcal{L}', w_1')$. If $\mathfrak{M} \models \langle a_1 \rangle ... \langle a_n \rangle !X_1$ and $\mathfrak{M}' \models \langle a_1 \rangle ... \langle a_n \rangle !X_2$ then there exists sequences $w_1, ..., w_{n+1}$ and $w_1', ..., w_{n+1}'$ of states such that

- $w_1 \xrightarrow{a_1} ... \xrightarrow{a_n} w_{n+1}$.
- $w_1' \xrightarrow{a_1} ... \xrightarrow{a_n} w_{n+1}'$.
- $\lambda(w_{n+1}) \subseteq X_1$.
- $\lambda'(w_{n+1}') \subseteq X_2$.

Now from the definition of $\mathsf{char}()$:

$$\mathsf{char}((\mathcal{L}, w_{n_1})) \vdash !X_1 \qquad \mathsf{char}((\mathcal{L}', w'_{n_1})) \vdash !X_2$$

Now using [!Right 2]:

$$\mathsf{char}((\mathcal{L}, w_{n_1})) \wedge \mathsf{char}((\mathcal{L}', w'_{n_1})) \vdash !(X_1 \cap X_2)$$

Using [!Right 1]:

$$\mathsf{char}((\mathcal{L}, w_{n_1})) \wedge \mathsf{char}((\mathcal{L}', w'_{n_1})) \vdash !X$$

Using n applications of [Normal]:

$$\langle a_1\rangle ... \langle a_n\rangle(\mathsf{char}((\mathcal{L}, w_{n_1})) \wedge \mathsf{char}((\mathcal{L}', w'_{n_1}))) \vdash \langle a_1\rangle ... \langle a_n\rangle !X$$

Finally, using n applications of [Det]:

$$\mathsf{char}((\mathcal{L}, w_1)) \wedge \mathsf{char}((\mathcal{L}', w'_1)) \vdash \langle a_1\rangle ... \langle a_n\rangle(\mathsf{char}((\mathcal{L}, w_{n_1})) \wedge \mathsf{char}((\mathcal{L}', w'_{n_1})))$$

So, by [Trans]

$$\mathsf{char}(\mathfrak{M}) \wedge \mathsf{char}(\mathfrak{M}') \vdash \langle a_1\rangle ... \langle a_n\rangle !X$$

□

8 Compactness and the standard translation to first-order logic

This section studies two embeddings of cathoristic logic into first-order logic. The second embedding is used to prove that cathoristic logic satisfies compactness.

8.1 Translating from cathoristic to first-order logic

The study of how a logic embeds into other logics is interesting in parts because it casts a new light on the logic that is the target of the embedding. A good example is the standard translation of modal into first-order logic. The translation produces various fragments: the finite variable fragments, the fragment closed under bisimulation, guarded fragments. These fragments have been investigated deeply, and found to have unusual properties not shared by the whole of first-order logic. Translations also enable us to push techniques, constructions and results between logics. In this section, we translate cathoristic logic into first-order logic.

Definition 14 The first-order signature $\mathcal{S}$ has a nullary predicate $\top$, a family of unary predicates $\mathsf{Restrict}_A(\cdot)$, one for each finite subset $A \subseteq \Sigma$, and a family of binary predicates $\mathsf{Arrow}_a(x, y)$, one for each action $a \in \Sigma$.

The intended interpretation is as follows.

- The universe is composed of states.
- The predicate $\top$ is true everywhere.
- For each finite $A \subseteq \Sigma$ and each state s, $\mathsf{Restrict}_A(s)$ is true if $\lambda(x) \subseteq A$.
- A set of two-place predicates $\mathsf{Arrow}_a(x, y)$, one for each $a \in \Sigma$, where x and y range over states. $\mathsf{Arrow}_a(x, y)$ is true if $x \xrightarrow{a} y$.

If Σ is infinite, then $\mathsf{Restrict}_A(\cdot)$ and $\mathsf{Arrow}_a(\cdot, \cdot)$ are infinite families of relations.

Definition 15 Choose two fixed variables x, y, let a range over actions in Σ, and A over finite subsets of Σ. Then the restricted fragment of first-order logic that is the target of our translation is given by the following grammar, where w, z range over x, y.

$$\phi \quad ::= \quad \top \mid \mathsf{Arrow}_a(w, z) \mid \mathsf{Restrict}_A(z) \mid \phi \wedge \psi \mid \exists x.\phi$$

This fragment has no negation, disjunction, implication, or universal quantification.

Definition 16 The translations $[\![\phi]\!]_x$ and $[\![\phi]\!]_y$ of cathoristic formula ϕ are given relative to a state, denoted by either x or y.

$$\begin{array}{rcl@{\qquad}rcl}
[\![\top]\!]_x & = & \top & [\![\top]\!]_y & = & \top \\
[\![\phi \wedge \psi]\!]_x & = & [\![\phi]\!]_x \wedge [\![\psi]\!]_x & [\![\phi \wedge \psi]\!]_y & = & [\![\phi]\!]_y \wedge [\![\psi]\!]_y \\
[\![\langle a \rangle \phi]\!]_x & = & \exists y.(\mathsf{Arrow}_a(x,y) \wedge [\![\phi]\!]_y) & [\![\langle a \rangle \phi]\!]_y & = & \exists x.(\mathsf{Arrow}_a(y,x) \wedge [\![\phi]\!]_x) \\
[\![!A]\!]_x & = & \mathsf{Restrict}_A(x) & [\![!A]\!]_y & = & \mathsf{Restrict}_A(y)
\end{array}$$

The translations on the left and right are identical, except for switching x and y. Here is an example translation.

$$[\![\langle a \rangle \top \wedge !\{a\}]\!]_x = \exists y.(\mathsf{Arrow}_a(x,y) \wedge \top) \wedge \mathsf{Restrict}_{\{a\}}(x)$$

We now establish the correctness of the encoding. The key issue is that not every first-order model of our first-order signature corresponds to a cathoristic model because determinism, well-sizedness and admissibility are not enforced by our signature alone. In other words, models may contain 'junk'. We deal with this problem following ideas from modal logic [5]: we add a translation $[\![\mathcal{L}]\!]$ for cathoristic transition systems, and then prove the following theorem.

Theorem 6 (correspondence theorem) *Let ϕ be a cathoristic logic formula and $\mathfrak{M} = (\mathcal{L}, s)$ a cathoristic model.*

$$\mathfrak{M} \models \phi \quad \textit{iff} \quad [\![\mathcal{L}]\!] \models_{x \mapsto s} [\![\phi]\!]_x.$$

And likewise for $[\![\phi]\!]_y$.

The definition of $[\![\mathcal{L}]\!]$ is simple.

Definition 17 Let $\mathcal{L} = (S, \to, \lambda)$ be a cathoristic transition system. Clearly $\mathcal{L}$ gives rise to an $\mathcal{S}$-model $[\![\mathcal{L}]\!]$ as follows.

- The universe is the set S of states.
- The relation symbols are interpreted as follows.
  - $\top^{[\![\mathcal{L}]\!]}$ always holds.
  - $\mathsf{Restrict}_A^{[\![\mathcal{L}]\!]} = \{s \in S \mid \lambda(s) \subseteq A\}$.
  - $\mathsf{Arrow}^{[\![\mathcal{L}]\!]}{}_a = \{(s,t) \in S \times S \mid s \xrightarrow{a} t\}$.

We are now ready to prove Theorem 6.

Proof: By induction on the structure of ϕ. The cases $\top$ and $\phi_1 \wedge \phi_2$ are straightforward. The case $\langle a \rangle \psi$ is handled as follows.

$$\begin{array}{ll}
[\![\mathcal{L}]\!] \models_{x \mapsto s} [\![\langle a \rangle \psi]\!]_x & \\
\quad \text{iff} & [\![\mathcal{L}]\!] \models_{x \mapsto s} \exists y.(\mathsf{Arrow}_a(x,y) \wedge [\![\psi]\!]_y) \\
\quad \text{iff} & \text{exists } t \in S. [\![\mathcal{L}]\!] \models_{x \mapsto s, y \mapsto t} \mathsf{Arrow}_a(x,y) \wedge [\![\psi]\!]_y \\
\quad \text{iff} & \text{exists } t \in S. [\![\mathcal{L}]\!] \models_{x \mapsto s, y \mapsto t} \mathsf{Arrow}_a(x,y) \text{ and } [\![\mathcal{L}]\!] \models_{x \mapsto s, y \mapsto t} [\![\psi]\!]_y \\
\quad \text{iff} & \text{exists } t \in S. s \xrightarrow{a} t \text{ and } [\![\mathcal{L}]\!] \models_{x \mapsto s, y \mapsto t} [\![\psi]\!]_y \\
\quad \text{iff} & \text{exists } t \in S. s \xrightarrow{a} t \text{ and } [\![\mathcal{L}]\!] \models_{y \mapsto t} [\![\psi]\!]_y \qquad (\text{as } x \text{ is not free in } \psi) \\
\quad \text{iff} & \text{exists } t \in S. s \xrightarrow{a} t \text{ and } \mathfrak{M} \models \psi \\
\quad \text{iff} & \mathfrak{M} \models \langle a \rangle \psi
\end{array}$$

Finally, if ϕ is $!A$ the derivation comes straight from the definitions.

$$\begin{array}{lll} [\![\mathcal{L}]\!] \models_{x \mapsto s} [\![!A]\!]_x & \text{iff} & [\![\mathcal{L}]\!] \models_{x \mapsto s} \mathsf{Restrict}_A(x) \\ & \text{iff} & \lambda(s) \subseteq A \\ & \text{iff} & \mathfrak{M} \models !A. \end{array}$$

□

8.2 Compactness by translation

First-order logic satisfies *compactness*: a set S of sentences has a model exactly when every finite subset of S does. What about cathoristic logic?

We can prove compactness of modal logics using the standard translation from modal to first-order logic [5]: we start from a set of modal formula such that each finite subset has a model. We translate the modal formulae and models to first-order logic, getting a set of first-order formulae such that each finite subset has a first-order model. By compactness of first-order logic, we obtain a first-order model of the translated modal formulae. Then we translate that first-order model back to modal logic, obtaining a model for the original modal formulae, as required. The last step proceeds without a hitch because the modal and the first-order notions of model are identical, save for details of presentation.

Unfortunately we cannot do the same with the translation from cathoristic logic to first-order logic presented in the previous section. The problem are the first-order models termed 'junk' above. The target language of the translation is not expressive enough to have formulae that can guarantee such constraints. As we have no reason to believe that the first-order model whose existence is guaranteed by compactness isn't 'junk', we cannot prove compactness with the translation. We solve this problem with a second translation, this time into a more expressive first-order fragment where we can constrain first-order models easily using formulae. The fragment we use now lives in two-sorted first-order logic (which can easily be reduced to first-order logic [11]).

Definition 18 The two-sorted first-order signature $\mathcal{S}'$ is given as follows.

- $\mathcal{S}'$ has two sorts, states and actions.
- The action constants are given by Σ. There are no state constants.
- $\mathcal{S}'$ has a nullary predicate $\top$.
- A binary predicate $\mathsf{Allow}(\cdot, \cdot)$. The intended meaning of $\mathsf{Allow}(x, a)$ is that at the state denoted by x we are allowed to do the action a.
- A ternary predicate $\mathsf{Arrow}(\cdot, \cdot, \cdot)$ where $\mathsf{Arrow}(x, a, y)$ means that there is a transition from the state denoted by x to the state denoted by y, and that transition is labelled a.

Definition 19 The encoding $\langle\!\langle \phi \rangle\!\rangle_x$ of cathoristic logic formulae is given by the following clauses.

$$\begin{aligned}
\langle\!\langle \top \rangle\!\rangle_x &= \top \\
\langle\!\langle \phi \wedge \psi \rangle\!\rangle_x &= \langle\!\langle \phi \rangle\!\rangle_x \wedge \langle\!\langle \psi \rangle\!\rangle_x \\
\langle\!\langle \langle a \rangle \phi \rangle\!\rangle_x &= \exists^{st} y.(\mathsf{Arrow}(x, a, y) \wedge \langle\!\langle \phi \rangle\!\rangle_y) \\
\langle\!\langle !A \rangle\!\rangle_x &= \forall^{act} a.(\mathsf{Allow}(x, a) \to a \in A)
\end{aligned}$$

Here we use $\exists^{st}$ to indicate that this existential quantifier ranges over the sort of states, and $\forall^{act}$ for the universal quantifier ranging over actions. The expression $a \in A$ is a shorthand for the first-order formula

$$a = a_1 \vee a = a_2 \vee \cdots \vee a = a_n$$

assuming that $A = \{a_1, ..., a_n\}$. Since by definition, A is always a finite set, this is well-defined. The translation could be restricted to a two-variable fragment. Moreover, the standard reduction from many-sorted to one-sorted first-order logic does not increase the number of variables used (although predicates are added, one per sort). We will not consider this matter further here.

We also translate cathoristic transition systems $\langle\!\langle \mathcal{L} \rangle\!\rangle$.

Definition 20 Let $\mathcal{L} = (S, \to, \lambda)$ be a cathoristic transition system. $\mathcal{L}$ gives rise to an $\mathcal{S}'$-model $\langle\!\langle \mathcal{L} \rangle\!\rangle$ as follows.

- The sort of states is interpreted by the set S.
- The sort of actions is interpreted by the set Σ.
- For each constant $a \in \Sigma$, $a^{\langle\!\langle \mathcal{L} \rangle\!\rangle}$ is a itself.
- The relation symbols are interpreted as follows.
  - $\top^{\langle\!\langle \mathcal{L} \rangle\!\rangle}$ always holds.
  - $\mathsf{Allow}^{\langle\!\langle \mathcal{L} \rangle\!\rangle}(s, a)$ holds whenever $a \in \lambda(s)$.
  - $\mathsf{Arrow}^{\langle\!\langle \mathcal{L} \rangle\!\rangle}(s, a, t)$ holds whenever $s \xrightarrow{a} t$.

Theorem 7 (correspondence theorem) *Let ϕ be a cathoristic logic formula and $\mathfrak{M} = (\mathcal{L}, s)$ a cathoristic model.*

$$\mathfrak{M} \models \phi \quad \textit{iff} \quad \langle\!\langle \mathcal{L} \rangle\!\rangle \models_{x \mapsto s} \langle\!\langle \phi \rangle\!\rangle_x.$$

Proof: The proof proceeds by induction on the structure of ϕ and is similar to that of Theorem 7. The case for the may modality proceeds as follows.

$$\begin{aligned}
\mathfrak{M} \models \langle a \rangle \phi \quad \text{iff} \quad & \text{exists state } t \text{ with } s \xrightarrow{a} t \text{ and } (\mathcal{L}, t) \models \phi \\
\text{iff} \quad & \text{exists state } t \text{ with } s \xrightarrow{a} t \text{ and } \langle\!\langle \mathcal{L} \rangle\!\rangle \models_{y \mapsto t} \langle\!\langle \phi \rangle\!\rangle_y \qquad \text{by (IH)} \\
\text{iff} \quad & \langle\!\langle \mathcal{L} \rangle\!\rangle \models_{x \mapsto s} \exists^{st} y.(\mathsf{Arrow}(x, a, y) \wedge \langle\!\langle \phi \rangle\!\rangle_y) \\
\text{iff} \quad & \langle\!\langle \mathcal{L} \rangle\!\rangle \models_{x \mapsto s} \langle\!\langle \langle a \rangle \phi \rangle\!\rangle_x
\end{aligned}$$

Finally $!A$.

$$\begin{array}{lll} \mathfrak{M} \models !A & \text{iff} & \lambda(s) \subseteq A \\ & \text{iff} & \text{for all } \ a \in \Sigma . a \in A \\ & \text{iff} & \langle\!\langle \mathcal{L} \rangle\!\rangle \models_{x \mapsto s} \forall^{act} a.(\mathsf{Allow}(x, a) \rightarrow a \in A) \\ & \text{iff} & \langle\!\langle \mathcal{L} \rangle\!\rangle \models_{x \mapsto s} \langle\!\langle !A \rangle\!\rangle_x \end{array}$$

□

We use the following steps in our compactness proof.

1. Choose a set Γ of cathoristic logic formulae such that each finite subset Γ' of Γ has a cathoristic model $(\mathcal{L}, s)$.
2. The translation gives a set $\langle\!\langle \Gamma \rangle\!\rangle = \{\langle\!\langle \phi \rangle\!\rangle \mid \phi \in \Gamma\}$ of first-order formulae such that each finite subset has a first-order model $\langle\!\langle \mathcal{L} \rangle\!\rangle$.
3. By compactness of (two-sorted) first-order logic, we can find a first-order model $\mathcal{M}$ of $\langle\!\langle \Gamma \rangle\!\rangle$.
4. Convert $\mathcal{M}$ into a cathoristic transition system $\mathcal{M}^\sharp$ such that $(\mathcal{M}^\sharp, s) \models \Gamma$.

The problematic step is (4) - for how would we know that the first-order model $\mathcal{M}$ can be converted back to a cathoristic transition system? What if it contains 'junk' in the sense described above? We solve this by adding formulae to $\langle\!\langle \Gamma \rangle\!\rangle$ that preserve finite satisfiability but force the first-order models to be convertible to cathoristic models. To ensure admissibility we use this formula.

$$\phi_{admis} \ = \ \forall^{st} s. \forall^{act} a. \forall^{st} t. (\mathsf{Arrow}(s, a, t) \rightarrow \mathsf{Allow}(s, a))$$

The formula ϕ_{det} ensures model determinism.

$$\phi_{det} \ = \ \forall^{st} s. \forall^{act} a. \forall^{st} t. \forall^{st} t'. ((\mathsf{Arrow}(s, a, t) \wedge \mathsf{Arrow}(s, a, t')) \rightarrow t = t')$$

Lemma 11 *If $\mathcal{L}$ is a cathoristic transition system then $\langle\!\langle \mathcal{L} \rangle\!\rangle \models \phi_{admis} \wedge \phi_{det}$.*

Proof: Straightforward from the definitions. □

We can now add, without changing satisfiability, $\phi_{admis} \wedge \phi_{det}$ to any set of first-order formulae that has a model that is the translation of a cathoristic model.

We also need to deal with well-sizedness in first-order models, because nothing discussed so far prevents models whose state labels are infinite sets without being Σ. Moreover, a model may interpret the set of actions with a proper superset of Σ. This also prevents conversion to cathoristic models. We solve these problems by simply removing all actions that are not in Σ and all transitions involving such actions. We map all infinite state labels to Σ. It is easy to see that this does not change satisfiability of (translations of) cathoristic formulae.

Definition 21 Let $\mathcal{L} = (S, \to, \lambda)$ be a cathoristic transition system and X a set, containing actions. The *restriction of $\mathcal{L}$ to X*, written $\mathcal{L} \setminus X$ is the cathoristic model $(S, \to', \lambda')$ where $\to' = \{(s, a, t) \in \to \mid a \notin X\}$, and for all states s we set:

$$\lambda'(s) = \begin{cases} \lambda(s) \setminus X & \text{whenever } \lambda(s) \neq \Sigma \\ \Sigma & \text{otherwise} \end{cases}$$

Lemma 12 *Let ϕ be a cathoristic logic formula and X be a set such that no action occurring in ϕ is in X. Then:*

$$(\mathcal{L}, s) \models \phi \quad \textit{iff} \quad (\mathcal{L} \setminus X, s) \models \phi.$$

Proof: By straightforward induction on the structure of ϕ, using the fact that by assumption X only contains actions not occurring in ϕ. □

Definition 22 Let $\mathcal{M}$ be a first-order model for the signature $\mathcal{S}'$. We construct a cathoristic transition system $\mathcal{M}^\sharp = (S, \to, \lambda)$.

- The actions Σ are given by the $\mathcal{M}$ interpretation of actions.
- The states S are given by the $\mathcal{M}$ interpretation of states.
- The reduction relation $s \xrightarrow{a} t$ holds exactly when $\mathsf{Arrow}^{\mathcal{M}}(s, a, t)$.
- The function λ is given by the following clause:

$$\lambda(s) = \begin{cases} X & \text{whenever } X = \{a \mid \mathsf{Allow}^{\mathcal{M}}(s, a)\} \text{ is finite} \\ \Sigma & \text{otherwise} \end{cases}$$

Lemma 13 *If $\mathcal{M}$ be a first-order model for $\mathcal{S}'$ such that $\mathcal{M} \models \phi_{admis} \wedge \phi_{det}$. Then $\mathcal{M}^\sharp$ is an cathoristic transition system with actions Σ.*

Proof: Immediate from the definitions. □

Theorem 8 (correspondence theorem) *Let $\mathcal{M}$ be a first-order model for the signature $\mathcal{S}'$ such that $\mathcal{M} \models \phi_{admis} \wedge \phi_{det}$. Then we have for all cathoristic logic formulae ϕ with actions from Σ:*

$$\mathcal{M} \models_{x \mapsto s} \langle\!\langle \phi \rangle\!\rangle_x \quad \textit{iff} \quad (\mathcal{M}^\sharp \setminus X, s) \models \phi.$$

Here X is the set of all elements in the universe of $\mathcal{M}$ interpreting actions that are not in Σ.

Proof: The proof proceeds by induction on the structure of ϕ. □

Definition 23 Let Γ be a set of cathoristic formulae, and $\mathfrak{M}$ a cathoristic model. We write $\mathfrak{M} \models T$ provided $\mathfrak{M} \models \phi$ for all $\phi \in T$. We say Γ is *satisfiable* provided $\mathfrak{M} \models T$.

Theorem 9 (Compactness of cathoristic logic) *A set Γ of cathoristic logic formulae is satisfiable iff each finite subset of Γ is satisfiable.*

Proof: For the non-trivial direction, let Γ be a set of cathoristic logic formulae such that any finite subset has a cathoristic model. Define

$$\langle\!\langle \Gamma \rangle\!\rangle \;=\; \{\langle\!\langle \phi \rangle\!\rangle \mid \phi \in \Gamma\} \qquad\qquad \Gamma^* \;=\; \langle\!\langle \Gamma \rangle\!\rangle \cup \{\phi_{admis} \wedge \phi_{det}\}$$

which both are sets of first-order formulae. Clearly each finite subset Γ' of Γ^* has a first-order model. Why? First consider the subset Γ'_{CL} of Γ' which is given as follows.

$$\Gamma'_{CL} \;=\; \{\phi \in \Gamma \mid \langle\!\langle \phi \rangle\!\rangle \in \Gamma'\}$$

Since Γ'_{CL} is finite, by assumption there is a cathoristic model

$$(\mathcal{L}, s) \models \Gamma'_{CL}$$

which means we can apply Theorem 8 to get

$$\langle\!\langle \mathcal{L} \rangle\!\rangle \models_{x \mapsto s} \langle\!\langle \Gamma'_{CL} \rangle\!\rangle,$$

By construction $\Gamma' \setminus \langle\!\langle \Gamma'_{CL} \rangle\!\rangle \subseteq \{\phi_{admis} \wedge \phi_{det}\}$, so all we have to show for Γ' to have a model is that

$$\langle\!\langle \mathcal{L} \rangle\!\rangle \models_{x \mapsto s} \{\phi_{admis}\} \cup \{\phi_a \mid a \in \Sigma\},$$

but that is a direct consequence of Lemma 11. That means each finite subset of Γ^* has a model and by appealing to compactness of first-order many-sorted logic (which is an immediate consequence of compactness of one-sorted first-order logic [11]), we know there must be a first-order model $\mathcal{M}$ of Γ^*, i.e.

$$\mathcal{M} \models \Gamma^*.$$

Since $\mathcal{M} \models \phi_{admis} \wedge \phi_{det}$ we can apply Theorem 8 that also

$$(\mathcal{M}^\sharp \setminus X, s) \models \Gamma$$

where X is the set of all actions in $\mathcal{M}^\sharp$ that are not in Σ. Hence Γ is satisfiable. □

9 Cathoristic logic and negation

We have presented cathoristic logic as a language that can express incompatible claims without negation. In this section, we briefly consider cathoristic logic enriched with negation.

9.1 Syntax and semantics

Definition 24 Given a set Σ of actions, the *formulae of cathoristic logic with negation* are given by the following grammar.

$$\phi \quad ::= \quad ... \mid \neg\phi$$

We can now define disjunction $\phi \vee \psi$ and implication $\phi \to \psi$ by de Morgan duality: $\phi \vee \psi$ is short for $\neg(\neg\phi \wedge \neg\psi)$, and $\phi \to \psi$ abbreviates $\neg\phi \vee \psi$.

The semantics of cathoristic logic with negation is just that of plain cathoristic logic except for the obvious clause for negation.

$$\mathfrak{M} \models \neg\phi \qquad \text{iff} \qquad \mathfrak{M} \not\models \phi$$

Negation is a core operation of classical logic, and its absence makes cathoristic logic unusual. In order to understand cathoristic logic better, we now investigate how negation can be seen as a definable abbreviation in cathoristic logic with disjunction. The key idea is to use the fact that

$$\neg\langle a\rangle\phi$$

can be false in two ways: either there is no a-labelled action at the current state - or there is, but ϕ is false. Both arms of this disjunction can be expressed in cathoristic logic, the former as $!\Sigma \setminus \{a\}$, the latter as $\langle a\rangle\neg\phi$. Hence, we can see $\neg\langle a\rangle\phi$ as a shorthand for

$$!(\Sigma \setminus \{a\}) \vee \langle a\rangle\neg\phi$$

Negation still occurs in this term, but prefixing a formula of lower complexity.

This leaves the question of negating the tantum. That's easy: when $\neg!A$, then clearly the current state can do an action $a \notin A$. In other words

$$\bigvee_{a\in\Sigma} \langle a\rangle\top$$

When Σ is infinite, then so is the disjunction.

Note that both the negation of the modality and the negation of the tantum involve the set Σ of actions. So far, we have defined negation with respect to the whole (possibly infinite) set Σ. For technical reasons, we generalise negation and define it with respect to a *finite* subset $S \subseteq \Sigma$. We use this finitely-restricted version of negation in the decision procedure below.

Definition 25 The function $\neg_S(\phi)$ removes negation from ϕ relative to a finite subset $S \subseteq \Sigma$:

$$\begin{array}{rclcrcl} \neg_S(\top) & = & \bot & & \neg_S(\bot) & = & \top \\ \neg_S(\phi \wedge \psi) & = & \neg_S(\phi) \vee \neg_S(\psi) & & \neg_S(\phi \vee \psi) & = & \neg_S(\phi) \wedge \neg_S(\psi) \\ \neg_S(\langle a \rangle \phi) & = & !(S - \{a\}) \vee \langle a \rangle \neg_S(\phi) & & \neg_S(!A) & = & \bigvee_{a \in S - A} \langle a \rangle \top \end{array}$$

9.2 Decision procedure

We can use the fact that cathoristic logic has a quadratic-time decision procedure to build a super-polynomial time decision procedure for cathoristic logic with negation. Given $\phi \models \psi$, let $S = \mathsf{actions}(\phi) \cup \mathsf{actions}(\psi) \cup \{a\}$, where a is a fresh action. The function $\mathsf{actions}(\cdot)$ returns all actions occurring in a formula, e.g. $\mathsf{actions}(\langle a \rangle \phi) = \{a\} \cup \mathsf{actions}(\phi)$ and $\mathsf{actions}(!A) = A$. The decision procedure executes the following steps.

1. Inductively translate away all negations in ϕ using $\neg_S(\phi)$ as defined above. Let the result be ϕ'.
2. Reduce ϕ' to disjunctive normal form by repeated application of the rewrite rules:

$$\phi \wedge (\psi \vee \xi) \rightsquigarrow (\phi \wedge \psi) \vee (\phi \wedge \xi) \qquad (\phi \vee \psi) \wedge \xi \rightsquigarrow (\phi \wedge \xi) \vee (\psi \wedge \xi).$$

3. Let the resulting disjuncts be $\phi_1, ..., \phi_n$. Note that

$$\phi \models \psi \quad \text{iff} \quad \phi_i \models \psi \text{ for all } i = 1, ..., n.$$

   For each disjunct ϕ_i do the following.
   - Notice that $\phi_i \models \psi$ if and only if all S-extensions (defined below) of $\mathsf{simpl}(\phi_i)$ satisfy ψ. So, to check whether $\phi_i \models \psi$, we enumerate the S-extensions of $\mathsf{simpl}(\phi_i)$ (there are a finite number of such extensions - the exact number is exponential in the size of $\mathsf{simpl}(\phi_i)$) and check for each such S-extension $\mathfrak{M}$ whether $\mathfrak{M} \models \psi$, using the algorithm of Section 6.5.

Here is the definition of S-extension.

Definition 26 Given an cathoristic transition system $\mathcal{L} = (\mathcal{W}, \rightarrow, \lambda)$, and a set S of actions, then $(\mathcal{W}', \rightarrow', \lambda')$ is a *S-extension* of $\mathcal{L}$ if it is a valid cathoristic transition system (recall Definition 2) and for all $(x, a, y) \in \rightarrow'$, either:

- $(x, a, y) \in \rightarrow$, or;
- $x \in \mathcal{W}$, $a \in S$, $a \in \lambda(x)$, and y is a new state not appearing elsewhere in $\mathcal{W}$ or $\mathcal{W}'$.

The state-labelling λ' is:

$$\lambda'(x) = \lambda(x) \text{ if } x \in \mathcal{W}$$
$$\lambda'(x) = \Sigma \text{ if } x \notin \mathcal{W}$$

In other words, $\mathfrak{M}'$ is an extension of an annotated model $\mathfrak{M}$, if all its transitions are either from $\mathfrak{M}$ or involve states of $\mathfrak{M}$ transitioning via elements of S to new states not appearing in $\mathfrak{M}$ or $\mathfrak{M}'$. The number of extensions grows quickly. If the model $\mathfrak{M}$ has n states, then the number of possible extensions is:

$$(2^{|S|})^n$$

But recall that we are computing these extensions in order to verify ψ. So we can make a significant optimisation by restricting the height of each tree to $|\psi|$. We state, without proof, that this optimisation preserves correctness. A Haskell implementation of the decision procedure is available [1].

10 Quantified cathoristic logic

So far, we have presented cathoristic logic as a propositional modal logic. This section sketches quantified cathoristic logic, primarily to demonstrate that this extension works smoothly.

Definition 27 Let Σ be a non-empty set of actions, ranged over by $a, a', ...$ as before. Given a set $\mathcal{V}$ of *variables*, with $x, x', y, y', ...$ ranging over $\mathcal{V}$, the *terms*, ranged over by $t, t', ...$ and formulae of quantified cathoristic logic are given by the following grammar:

$$\begin{array}{lcl} t & ::= & x \mid a \\ \phi & ::= & \top \mid \phi \wedge \psi \mid \langle t \rangle \phi \mid !A \mid \exists x.\phi \mid \forall x.\phi \end{array}$$

Now A ranges over finite subsets of terms. The *free variables* of a ϕ, denoted $\mathsf{fv}(\phi)$ is given as expected, e.g. $\mathsf{fv}(\langle t \rangle \phi) = \mathsf{fv}(t) \cup \mathsf{fv}(\phi)$ and $\mathsf{fv}(!A) = \bigcup_{t \in A} \mathsf{fv}(t)$ where $\mathsf{fv}(a) = \emptyset$ and $\mathsf{fv}(x) = \{x\}$.

Definition 28 The semantics of quantified cathoristic logic is constructed along conventional lines. An *environment* is a map $\sigma : \mathcal{V} \to \Sigma$ with finite domain. We write $\sigma, x : a$ for the environment that is just like σ, except it also maps x to a, implicitly assuming that x is not in σ's domain. The *denotation* $[\![t]\!]_\sigma$ of a term t under an environment σ is given as follows:

$$[\![a]\!]_\sigma = a \qquad\qquad [\![x]\!]_\sigma = \sigma(x)$$

where we assume that $\mathsf{fv}(t)$ is a subset of the domain of σ.

The *satisfaction relation* $\mathfrak{M} \models_\sigma \phi$ is defined whenever $\mathsf{fv}(\phi)$ is a subset of σ's domain. It is given by the following clauses, where we assume that $\mathfrak{M} = (\mathcal{L}, s)$ and $\mathcal{L} = (S, \to, \lambda)$.

$$\begin{array}{lll} \mathfrak{M} \models_\sigma \top & & \\ \mathfrak{M} \models_\sigma \phi \wedge \psi & \text{iff} & \mathfrak{M} \models_\sigma \phi \text{ and } \mathfrak{M} \models_\sigma \psi \\ \mathfrak{M} \models_\sigma \langle t \rangle \phi & \text{iff} & \text{there is transition } s \xrightarrow{[\![t]\!]_\sigma} s' \text{ such that } (\mathcal{L}, s') \models_\sigma \phi \\ \mathfrak{M} \models_\sigma !A & \text{iff} & \lambda(s) \subseteq \{[\![t]\!] \mid t \in A\} \\ \mathfrak{M} \models_\sigma \forall x.\phi & \text{iff} & \text{for all } a \in \Sigma \text{ we have } \mathfrak{M} \models_{\sigma, x:a} \phi \\ \mathfrak{M} \models_\sigma \exists x.\phi & \text{iff} & \text{there exists } a \in \Sigma \text{ such that } \mathfrak{M} \models_{\sigma, x:a} \phi \end{array}$$

In quantified cathoristic logic, we can say that there is exactly one king of France, and he is bald, as:

$$\exists x.(\langle king \rangle \langle france \rangle !\{x\} \wedge \langle x \rangle \langle bald \rangle)$$

Expressing this in first-order logic is more cumbersome:

$$\exists x.(king(france, x) \wedge bald(x) \wedge \forall y.(king(france, y) \to y = x))$$

The first-order logic version uses an extra universal quantifier, and also requires the identity relation with concomitant axioms.

To say that every person has exactly one sex, which is either male or female, we can write in quantified cathoristic logic:

$$\forall x.(\langle x\rangle\langle person\rangle \rightarrow \langle x\rangle\langle sex\rangle!\{male, female\} \wedge \exists y.\langle x\rangle\langle sex\rangle(\langle y\rangle\wedge!\{y\}))$$

This is more elegant than the equivalent in first-order logic:

$$\forall x.(person(x) \rightarrow \exists y. \begin{pmatrix} sex(x,y) \\ \quad \wedge \\ (y = male \ \vee \ y = female) \\ \quad \wedge \\ \forall z.sex(x,z) \rightarrow y = z \end{pmatrix})$$

To say that every traffic light is coloured either green, amber or red, we can write in quantified cathoristic logic:

$$\forall x.(\langle x\rangle\langle light\rangle \rightarrow \langle x\rangle\langle colour\rangle!\{green, amber, red\} \wedge \exists y.\langle x\rangle\langle colour\rangle(\langle y\rangle\wedge!\{y\}))$$

Again, this is less verbose than the equivalent in first-order logic:

$$\forall x.(light(x) \rightarrow \exists y. \begin{pmatrix} colour(x,y) \\ \quad \wedge \\ (y = green \ \vee \ y = amber \ \vee \ y = red) \\ \quad \wedge \\ \forall z.colour(x,z) \rightarrow y = z \end{pmatrix})$$

11 Related work

This section surveys cathoristic logic's intellectual background, and related approaches.

11.1 Brandom's incompatibility semantics

In [8], Chapter 5, Appendix I, Brandom developed a new type of semantics, incompatibility semantics, that takes material incompatibility - rather than truth-assignment - as the semantically primitive notion.

Incompatibility semantics applies to any language, $\mathcal{L}$, given as a set of sentences. Given a predicate $\mathsf{Inc}(X)$ which is true of sets $X \subseteq \mathcal{L}$ that are incompatible, he defines an incompatibility function $\mathcal{I}$ from subsets of $\mathcal{L}$ to sets of subsets of $\mathcal{L}$:

$$X \in \mathcal{I}(Y) \quad \text{iff} \quad \mathsf{Inc}(X \cup Y).$$

We assume that $\mathcal{I}$ satisfies the monotonicity requirement (Brandom calls it "Persistence"):

$$\text{If } X \in \mathcal{I}(Y) \text{ and } X \subseteq X' \text{ then } X' \in \mathcal{I}(Y).$$

Now Brandom defines entailment in terms of the incompatibility function. Given a set $X \subseteq \mathcal{L}$ and an individual sentence $\phi \in \mathcal{L}$:

$$X \models \phi \quad \text{iff} \quad \mathcal{I}(\{\phi\}) \subseteq \mathcal{I}(X).$$

Now, given material incompatibility (as captured by the $\mathcal{I}$ function) and entailment, he introduces logical negation as a *derived* concept via the rule:

$$\{\neg\phi\} \in \mathcal{I}(X) \quad \text{iff} \quad X \models \phi.$$

Brandom goes on to show that the $\neg$ operator, as defined, satisfies the laws of classical negation. He also introduces a modal operator, again defined in terms of material incompatibility, and shows that this operator satisfies the laws of $S5$.

Cathoristic logic was inspired by Brandom's vision that material incompatibility is conceptually prior to logical negation: in other words, it is possible for a community of language users to make incompatible claims, even if that language has no explicit logical operators such as negation. The language users of this simple language may go on to introduce logical operators, in order to make certain inferential properties explicit - but this is an optional further development. The language before that addition was already in order as it is.

The approach taken in this paper takes Brandom's original insight in a different direction. While Brandom defines an unusual (non truth-conditional) semantics that applies to any language, we have defined an unusual logic with a standard (truth-conditional) semantics, and then shown that this logic satisfies the Brandomian connection between incompatibility and entailment.

11.2 Peregrin on defining a negation operator

Peregrin [22] investigates the structural rules that any logic must satisfy if it is to connect incompatibility (Inc) and entailment ($\models$) via the Brandomian incompatibility semantics constraint:

$$X \models \phi \quad \text{iff} \quad \mathcal{I}(\{\phi\}) \subseteq \mathcal{I}(X).$$

The general structural rules are:

$$\begin{aligned}
(\bot)\ & \text{If } \mathsf{Inc}(X) \text{ and } X \subseteq Y \text{ then } \mathsf{Inc}(Y).\\
(\models 1)\ & \phi, X \models \phi.\\
(\models 2)\ & \text{If } X, \phi \models \psi \text{ and } Y \models \phi \text{ then } X, Y \models \psi.\\
(\bot \models 2)\ & \text{If } X \models \phi \text{ for all } \phi, \text{ then } \mathsf{Inc}(X).\\
(\models \bot 2)\ & \text{If } \mathsf{Inc}(Y \cup \{\phi\}) \text{ implies } \mathsf{Inc}(Y \cup X) \text{ for all } Y, \text{ then } X \models \phi.
\end{aligned}$$

Peregrin shows that if a logic satisfied the above laws, then incompatibility and entailment are mutually interdefinable, and the logic satisfies the Brandomian incompatibility semantics constraint.

Next, Peregrin gives a pair of laws for defining negation in terms of Inc and $\models$[16]:

$$\begin{aligned}
(\neg 1)\ & \mathsf{Inc}(\{\phi, \neg\phi\}).\\
(\neg 2)\ & \text{If } \mathsf{Inc}(X, \phi) \text{ then } X \models \neg\phi.
\end{aligned}$$

These laws characterise intuitionistic negation as the *minimal incompatible*[17]. Now, in [8], Brandom defines negation slightly differently. He uses the rule:

$$(\neg B)\ \ \mathsf{Inc}(X, \neg\phi) \text{ iff } X \models \phi.$$

Using this stronger rule, we can infer the classical law of double-negation: $\neg\neg\phi \models \phi$. Peregrin establishes that Brandom's rule for negation entail ($\neg 1$) and ($\neg 2$) above, but not conversely: Brandom's rule is stronger than Peregrin's minimal laws ($\neg 1$) and ($\neg 2$).

Peregrin concludes that the Brandomian constraint between incompatibility and entailment is satisfied by many different logics. Brandom happened to choose a particular rule for negation that led to classical logic, but the general connection between incompatibility and entailment is satisfied by many different logics, including intuitionistic logic. This paper supports Peregrin's conclusion: we have shown that cathoristic logic also satisfies the Brandomian constraint.

[16] The converse of ($\neg 2$) follows from ($\neg 1$) and the general structural laws above.

[17] ψ is the minimal incompatible of ϕ iff for all ξ, if $\mathsf{Inc}(\{\phi\} \cup \{\xi\})$ then $\xi \models \psi$.

11.3 Peregrin and Turbanti on defining a necessity operator

In [8], Brandom gives a rule for defining necessity in terms of incompatibility and entailment:

$$X \in \mathcal{I}(\{\Box\phi\}) \quad \text{iff} \quad \mathsf{Inc}(X) \vee \exists Y.\, Y \notin \mathcal{I}(X) \wedge Y \nvDash \phi.$$

In other words, X is incompatible with $\Box\phi$ if X is compatible with something that does not entail ϕ.

The trouble is, as Peregrin and Turbanti point out, if ϕ is not tautological, then *every set* $X \subseteq \mathcal{L}$ is incompatible with $\Box\phi$. To show this, take any set $X \subseteq \mathcal{L}$. If $\mathsf{Inc}(X)$, then $X \in \mathcal{I}(\Box\phi)$ by definition. If, on the other hand, $\neg\mathsf{Inc}(X)$, then let $Y = \emptyset$. Now $\neg\mathsf{Inc}(X \cup Y)$ as $Y = \emptyset$, and $Y \nvDash \phi$ as ϕ is not tautological. Hence $X \in \mathcal{I}(\Box\phi)$ for all $X \subseteq \mathcal{L}$. Brandom's rule, then, is only capable of specifying a very specific form of necessity: logical necessity.

In [22] and [31], Peregrin and Turbanti describe alternative ways of defining necessity. These alternative rule sets can be used to characterise modal logics other than S5. For example, Turbanti defines the accessibility relation between worlds in terms of a *compossibility relation*, and then argues that the S4 axiom of transitivity fails because compossibility is not transitive.

We draw two conclusions from this work. The first is, once again, that a commitment to connecting incompatibility and entailment via the Brandomian constraint:

$$X \models \phi \quad \text{iff} \quad \mathcal{I}(\{\phi\}) \subseteq \mathcal{I}(X)$$

does not commit us to any particular logical system. There are a variety of logics that can satisfy this constraint. Second, questions about the structure of the accessibility relation in Kripke semantics - questions that can seem hopelessly abstract and difficult to answer - can be re-cast in terms of concrete questions about the incompatibility relation. Incompatibility semantics can shed light on possible-world semantics [31].

11.4 Linear logic

Linear logic [15] is a refinement of first-order logic and was introduced by J.-Y. Girard and brings the symmetries of classical logic to constructive logic.

Linear logic splits conjunction into additive and multiplicative parts. The former, additive conjunction $A\&B$, is especially interesting in the context of cathoristic logic. In the terminology of process calculus it can be interpreted as an external choice operation [2]. ('External', because the choice is offered to the environment). This interpretation has been influential in the study of types for process calculus, e.g. [19, 20, 29]. Implicitly, additive conjunction gives an explicit upper bound on how many different options the environment can choose from. For example $A\&B\&C$ has three options (assuming that none of A, B, C can be decomposed into further additive conjunctions). With this in mind, and simplifying a great deal, a key difference between $!A$ and additive

conjunction $A\&B$ is that the individual actions in $!A$ have no continuation, while they do with $A\&B$: the tantum $!\{l, r\}$ says that the only permitted actions are l and r. What happens at later states is not constrained by $!A$. In contrast, $A\&B$ says not only that at this point the only permissible options are A and B, but also that if we choose A, then A holds 'for ever', and likewise for choosing B. To be sure, the alternatives in $A\&B$ may themselves contain further additive conjunctions, and in this way express how exclusion changes 'over time'.

In summary, cathoristic logic and linear logic offer operators that restrict the permissible options. How are they related? Linear logic has an explicit linear negation $(\cdot)^{\perp}$ which, unlike classical negation, is constructive. In contrast, cathoristic logic defines a restricted form of negation using $!A$. Can these two perspectives be fruitfully reconciled?

11.5 Process calculus

Process calculi are models of concurrent computation. They are based on the idea of message passing between actors running in parallel. Labelled transition systems are often used as models for process calculi, and many concepts used in the development of cathoristic logic - for example, bisimulations and Hennessy-Milner logic - originated in process theory (although some, such as bisimulation, evolved independently in other contexts).

Process calculi typically feature a construct called sum, that is an explicit description of mutually exclusive option:

$$\sum_{i \in I} P_i$$

That is a process that can internally choose, or be chosen externally by the environment to evolve into the process P_i for each i. Once the choice is made, all other options disappear. Sums also relate closely to linear logic's additive conjunction. Is this conceptual proximity a coincidence or indicative of deeper common structure?

11.6 Linguistics

Linguists have also investigated how mutually exclusive alternatives are expressed, often in the context of antonymy [3,4,21], but, to the best of our knowledge have not proposed formal theories of linguistic exclusion.

12 Open problems

In this paper, we have introduced cathoristic logic and established key meta-logical properties. However, many questions are left open.

12.1 Excluded middle

One area we would like to investigate further is what happens to the law of excluded middle in cathoristic logic. The logical law of excluded middle states that either a proposition or its negation must be true. In cathoristic logic

$$\models \phi \vee \neg_S(\phi)$$

does not hold in general. (The negation operator $\neg_S(\cdot)$ was defined in Section 9.) For example, let ϕ be $\langle a\rangle\top$ and $S = \Sigma = \{a, b\}$. Then

$$\phi \vee \neg_S\phi \quad = \quad \langle a\rangle\top \vee !\{b\} \ \vee \ \langle a\rangle\bot$$

Now this will not in general be valid - it will be false for example in the model $((\{x\}, \emptyset, \{(x, \Sigma)\}), x)$, the model having just the start state (labelled Σ) and no transitions. Restricting S to be a proper subset of $\Sigma = \{a, b\}$ is also not enough. For example with $S = \{a\}$ we have

$$\langle a\rangle\top \vee \neg_S(\langle a\rangle\top) \quad = \quad \langle a\rangle\top \vee !\emptyset \vee \langle a\rangle\bot$$

This formula cannot hold in any cathoristic model which contains a b-labelled transition, but no a-transition from the start state.

Is it possible to identify classes of models that nevertheless verify excluded middle? The answer to this question appears to depend on the chosen notion of semantic model.

12.2 Understanding the expressive strength of cathoristic logic

12.2.1 Comparing cathoristic logic and Hennessy-Milner logic

Section 8.1 investigated the relationship between cathoristic logic and first-order logic. Now we compare cathoristic logic with a logic that is much closer in spirit: Hennessy-Milner logic [17], a multi-modal logic designed to reason about process calculi. Indeed, the present shape of cathoristic logic owes much to Hennessy-Milner logic. We contrast both by translation from the former into the latter. This will reveal, more clearly than the translation into first-order logic, the novelty of cathoristic logic.

Definition 29 Assume a set Σ of symbols, with s ranging over Σ, the *formulae* of Hennessy-Milner logic are given by the following grammar:

$$\phi \quad ::= \quad \top \mid \textstyle\bigwedge_{i\in I} \phi_i \mid \langle s\rangle\phi \mid \neg\phi$$

The index set I in the conjunction can be infinite, and needs to be so for applications in process theory.

Definition 30 *Models* of Hennessy-Milner logic are simply pairs $(\mathcal{L}, s)$ where $\mathcal{L} = (S, \to)$ is a labelled transition system over Σ, and $s \in S$. The *satisfaction relation* $(\mathcal{L}, s) \models \phi$ is given by the following inductive clauses.

$$\begin{array}{lll} (\mathcal{L}, s) \models \top & & \\ (\mathcal{L}, s) \models \bigwedge_{i \in I} \phi_i & \text{iff} & \text{for all } i \in I \ : (\mathcal{L}, s) \models \phi_i \\ (\mathcal{L}, s) \models \langle a \rangle \phi & \text{iff} & \text{there is a } s \xrightarrow{a} s' \text{ such that } (l, s') \models \phi \\ (\mathcal{L}, s) \models \neg \phi & \text{iff} & (\mathcal{L}, s) \not\models \phi \end{array}$$

There are two differences between cathoristic logic and Hennessy-Milner logic - one syntactic, the other semantic.

- Syntactically, cathoristic logic has the tantum operator (!) instead of logical negation ($\neg$).
- Semantically, cathoristic models are deterministic, while (typically) models of Hennessy-Milner logic are non-deterministic (although the semantics makes perfect sense for deterministic transition systems, too). Moreover, models of Hennessy-Milner logic lack state labels.

Definition 31 We translate formulae of cathoristic logic into Hennessy-Milner logic using the function $[\![\cdot]\!]$:

$$\begin{aligned} [\![\top]\!] &= \top \\ [\![\phi_1 \wedge \phi_2]\!] &= [\![\phi_1]\!] \wedge [\![\phi_2]\!] \\ [\![\langle a \rangle \phi]\!] &= \langle a \rangle [\![\phi]\!] \\ [\![!A]\!] &= \bigwedge_{a \in \Sigma \setminus A} \neg \langle a \rangle \top \end{aligned}$$

If Σ is an infinite set, then the translation of a !-formula will be an infinitary conjunction. If Σ is finite, then the size of the Hennessy-Milner logic formula will be of the order of $n \cdot |\Sigma|$ larger than the original cathoristic formula, where n is the number of tantum operators occurring in the cathoristic formula). In both logics we use the number of logical operators as a measure of size.

We can also translate cathoristic models by forgetting state-labelling:

$$[\![((S, \to, \lambda), s)]\!] = ((S, \to), s)$$

We continue with an obvious consequence of the translation.

Theorem 10 *Let $\mathfrak{M}$ be a (deterministic or non-deterministic) cathoristic model. Then $\mathfrak{M} \models \phi$ implies $[\![\mathfrak{M}]\!] \models [\![\phi]\!]$.*

Proof: Straightforward by induction on ϕ. □

However, note that the following natural extension is *not* true under the translation above:

$$\text{If } \phi \models \psi \text{ then } [\![\phi]\!] \models [\![\psi]\!]$$

To see this, consider an entailment which relies on determinism, such as

$$\langle a\rangle\langle b\rangle \wedge \langle a\rangle\langle c\rangle \models \langle a\rangle(\langle b\rangle \wedge \langle c\rangle)$$

The first entailment is valid in cathoristic logic because of the restriction to deterministic models, but not in Hennessy-Milner logic, where it is invalidated by any model with two outgoing a transitions, one of which satisfies $\langle b\rangle$ and one of which satisfies $\langle c\rangle$.

We can restore the desired connection between cathoristic implication and Hennessy-Milner logic implication in two ways. First we can restrict our attention to deterministic models of Hennessy-Milner logic. The second solution is to add a determinism constraint to our translation. Given a set Γ of cathoristic formulae, closed under sub formulae, that contains actions from the set $A \subseteq \Sigma$, let the determinism constraint for Γ be:

$$\bigwedge_{a\in A,\phi\in\Gamma,\psi\in\Gamma} \neg\left(\langle a\rangle\phi \wedge \langle a\rangle\psi \wedge \neg\langle a\rangle(\phi \wedge \psi)\right)$$

If we add this sentence as part of our translation $[\![\cdot]\!]$, we do get the desired result that

$$\text{If } \phi \models \psi \text{ then } [\![\phi]\!] \models [\![\psi]\!]$$

12.2.2 Comparing cathoristic logic with Hennessy-Milner logic and propositional logic

Consider the following six languages:

Language	Description
PL[$\wedge$]	Propositional logic without negation
Hennessy-Milner logic[$\wedge$]	Hennessy-Milner logic without negation
CL[$\wedge, !$]	Cathoristic logic
PL [$\wedge, \neg$]	Full propositional logic
HML [$\wedge, \neg$]	Full Hennessy-Milner logic
CL [$\wedge, !, \neg$]	Cathoristic logic with negation

The top three languages are simple. In each case: there is no facility for expressing disjunction, every formula that is satisfiable has a simplest satisfying model, and there is a simple quadratic-time decision procedure But there are two ways in which CL[$\wedge, !$] is more expressive. Firstly, CL[$\wedge, !$], unlike HML[$\wedge$], is expressive enough to be able to distinguish between any two models that are not bisimilar, cf. Theorem 11. The second way in which CL[$\wedge, !$] is significantly more expressive than both PL[$\wedge$] and HML[$\wedge$] is in its ability to express incompatibility. No two formulae of PL[$\wedge$] or HML[$\wedge$] are incompatible[18] with each other. But many pairs of formulae of CL[$\wedge, !$] are incompatible.

[18] The notion of incompatibility applies to all logics: two formulae are incompatible if there is no model which satisfies both.

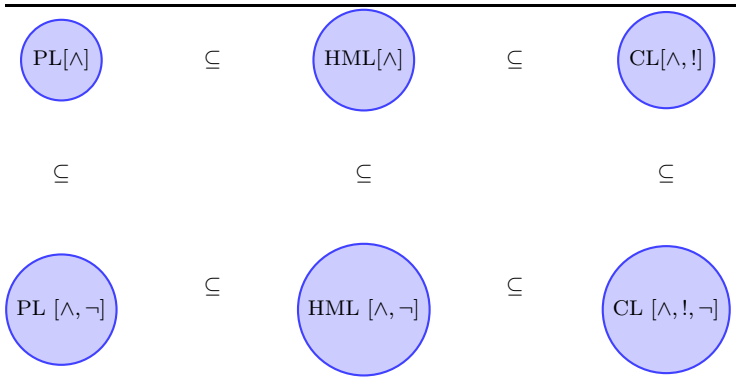

Fig. 15: Conjectured relationships of expressivity between logics. Here $L_1 \subseteq L_2$ means that the logic L_2 is more expressive than L_1. We leave the precise meaning of logical expressivity open.

(For example: $\langle a\rangle\top$ and $!\emptyset$). Because CL$[\wedge, !]$ is expressive enough to be able to make incompatible claims, it satisfies Brandom's incompatibility semantics constraint. CL$[\wedge, !]$ is the only logic (we are aware of) with a quadratic-time decision procedure that is expressive enough to respect this constraint.

The bottom three language can all be decided in super-polynomial time. We claim that Hennessy-Milner logic is more expressive than PL, and CL$[\wedge, !, \neg]$ is more expressive than full Hennessy-Milner logic. To see that full Hennessy-Milner logic is more expressive than full propositional logic, fix a propositional logic with the nullary operator $\top$ plus an infinite number of propositional atoms $P_{(i,j)}$, indexed by i and j. Now translate each formula of Hennessy-Milner logic via the rules:

$$[\![\top]\!] = \top \qquad [\![\phi \wedge \psi]\!] = [\![\phi]\!] \wedge [\![\psi]\!]$$
$$[\![\neg\phi]\!] = \neg[\![\phi]\!] \qquad [\![\langle a_i\rangle\phi_j]\!] = P_{(i,j)}$$

We claim Hennessy-Milner logic is more expressive because there are formulae ϕ and ψ of Hennessy-Milner logic such that

$$\phi \models_{\text{HML}} \psi \text{ but } [\![\phi]\!] \not\models_{\text{PL}} [\![\psi]\!]$$

For example, let $\phi = \langle a\rangle\langle b\rangle\top$ and $\psi = \langle a\rangle\top$. Clearly, $\phi \models_{\text{HML}} \psi$. But $[\![\phi]\!] = P_{(i,j)}$ and $[\![\psi]\!] = P_{(i',j')}$ for some i, j, i', j', and there are no entailments in propositional logic between arbitrary propositional atoms.

We close by stating that CL$[\wedge, !, \neg]$ is more expressive than full Hennessy-Milner logic. As mentioned above, the formula $!A$ of cathoristic logic can be translated into Hennessy-Milner logic as:

$$\bigwedge_{a\in\Sigma-A} \neg\langle a\rangle\top$$

But if Σ is infinite, then this is an infinitary disjunction. Cathoristic logic can express the same proposition in a finite sentence.

12.3 Acknowledgements

We thank Tom Smith and Giacomo Turbanti for their thoughtful comments.

References

1. Haskell implementation of cathoristic logic. Submitted with the paper, 2014.
2. S. Abramsky. Computational interpretations of linear logic. *TCS*, 111, 1993.
3. K. Allan, editor. *Concise Encyclopedia of Semantics*. Elsevier, 2009.
4. M. Aronoff and J. Rees-Miller, editors. *The Handbook of Linguistics*. Wiley-Blackwell, 2003.
5. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
6. R. Brachman and H. Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann, 2004.
7. R. Brandom. *Making It Explicit*. Harvard University Press, 1998.
8. R. Brandom. *Between Saying and Doing*. Oxford University Press, 2008.
9. B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, 1990.
10. D. Davidson. *Essays on Actions and Events*. Oxford University Press, 1980.
11. H. B. Enderton. *A Mathematical Introduction to Logic*. Academic Press, 2001.
12. R. Evans and E. Short. Versu. `http://www.versu.com`, available at `https://itunes.apple.com/us/app/blood-laurels/id882505676?mt=8`.
13. R. Evans and E. Short. Versu - a simulationist storytelling system. *IEEE Transactions on Computational Intelligence and AI in Games*, 2014.
14. R. Fikes and N. Nilsson. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2, 1971.
15. J.-Y. Girard. Linear logic. *TCS*, 50, 1987.
16. M. Hennessy. *Algebraic theory of processes*. MIT Press series in the foundations of computing. MIT Press, 1988.
17. M. Hennessy and R. Milner. Algebraic Laws for Non-Determinism and Concurrency. *JACM*, 32(1), 1985.
18. K. Honda. A Theory of Types for the π-Calculus. Available at: `http://www.dcs.qmul.ac.uk/~kohei/logics`, March 2001.
19. K. Honda, V. T. Vasconcelos, and M. Kubo. Language primitives and type disciplines for structured communication-based programming. In *Proc. ESOP*, volume 1381 of *LNCS*, pages 22–138, 1998.
20. K. Honda and N. Yoshida. A uniform type structure for secure information flow. *SIGPLAN Not.*, 37:81–92, January 2002.
21. A. O'Keeffe and M. McCarthy, editors. *The Routledge Handbook of Corpus Linguistics*. Routledge, 2010.
22. J. Peregrin. Logic as based on incompatibility. Available from `http://philpapers.org/rec/PERLAB-2`, 2010.
23. A. M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2013.
24. B. Russell. *An Inquiry into Meaning and Truth*. Norton and Co, 1940.
25. D. Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge University Press, 2012.
26. V. Sassone, M. Nielsen, and G. Winskel. Models for Concurrency: Towards a Classification. *TCS*, 170(1-2):297–348, 1996.
27. D. Smith and M. Genesereth. Ordering conjunctive queries. *Artificial Intelligence*, 26, 1985.
28. F. Sommers. *The Logic of Natural Language*. Clarendon Press, 1982.
29. K. Takeuchi, K. Honda, and M. Kubo. An Interaction-based Language and its Typing System. In *Proc. PARLE*, volume 817 of *LNCS*, pages 398–413, 1994.
30. H. Troelstra, A. S. and Schwichtenberg. *Basic proof theory (2nd ed.)*. Cambridge University Press, 2000.
31. G. Turbanti. Modality in Brandom's Incompatibility Semantics. In *Proceedings of the Amsterdam Graduate Conference - Truth, Meaning, and Normativity*, 2011.
32. D. van Dalen. *Logic and Structure*. Springer Verlag, 2004.
33. L. Wittgenstein. *Philosophische Bemerkungen*. Suhrkamp Verlag, 1981. Edited by R. Rhees.
34. L. Wittgenstein. *Tractatus logico-philosophicus: Logisch-philosophische Abhandlung*. Suhrkamp Verlag, 2003. Originally published: 1921.

A Alternative semantics for cathoristic logic

We use state-labelled transition systems as models for cathoristic logic. The purpose of the labels on states is to express constraints, if any, on outgoing actions. This concern is reflected in the semantics of $!A$.

$$((S, \to, \lambda), s) \models !A \quad \text{iff} \quad \lambda(s) \subseteq A$$

There is an alternative, and in some sense even simpler approach to giving semantics to $!A$ which does not require state-labelling: we simply check if all actions of all outgoing transitions at the current state are in A. As the semantics of other formula requires state-labelling in its satisfaction condition, this means we can use plain labelled transition systems (together with a current state) as models. This gives rise to a subtly different theory that we now explore, albeit not in depth.

A.1 Pure cathoristic models

Definition 32 By a *pure cathoristic model*, ranged over by $\mathfrak{P}, \mathfrak{P}', ...$, we mean a pair $(\mathcal{L}, s)$ where $\mathcal{L} = (S, \to)$ is a deterministic labelled transition system and $s \in S$ a state.

Adapting the satisfaction relation to pure cathoristic models is straightforward.

Definition 33 Using pure cathoristic models, the *satisfaction relation* is defined inductively by the following clauses, where we assume that $\mathfrak{M} = (\mathcal{L}, s)$ and $\mathcal{L} = (S, \to)$.

$$\begin{array}{lll} \mathfrak{M} \models \top & & \\ \mathfrak{M} \models \phi \wedge \psi & \text{iff} & \mathfrak{M} \models \phi \text{ and } \mathfrak{M} \models \psi \\ \mathfrak{M} \models \langle a \rangle \phi & \text{iff} & \text{there is a } s \xrightarrow{a} t \text{ such that } (\mathcal{L}, t) \models \phi \\ \mathfrak{M} \models A & \text{iff} & \{a \mid \exists t. s \xrightarrow{a} t\} \subseteq A \end{array}$$

Note that all but the last clause are unchanged from Definition 4.

In this interpretation, $!A$ restricts the out-degree of the current state s, i.e. it constraints the 'width' of the graph. It is easy to see that all rules in Figure 12 are sound with respect to the new semantics. The key advantage pure cathoristic models have is their simplicity: they are unadorned labelled transition systems, the key model of concurrency theory [26]. The connection with concurrency theory is even stronger than that, because, as we show below (Theorem 11), the elementary equivalence on (finitely branching) pure cathoristic models is bisimilarity, one of the more widely used notions of process equivalence. This characterisation even holds if we remove the determinacy restriction in Definition 32.

A.2 Relationship between pure and cathoristic models

The obvious way of converting an cathoristic model into a pure cathoristic model is by forgetting about the state-labelling:

$$((S, \to, \lambda), s) \qquad \mapsto \qquad ((S, \to), s)$$

Let this function be $\mathsf{forget}(\cdot)$. For going the other way, we have two obvious choices:

- $((S, \to), s) \mapsto ((S, \to, \lambda), s)$ where $\lambda(t) = \Sigma$ for all states t. Call this map $\mathsf{max}(\cdot)$.
- $((S, \to), s) \mapsto ((S, \to, \lambda), s)$ where $\lambda(t) = \{a \mid \exists t'. t \xrightarrow{a} t'\}$ for all states t. Call this map $\mathsf{min}(\cdot)$.

Lemma 14 *Let $\mathfrak{M}$ be an cathoristic model, and $\mathfrak{P}$ a pure cathoristic model.*

1. *$\mathfrak{M} \models \phi$ implies $\mathsf{forget}(\mathfrak{M}) \models \phi$. The reverse implication does not hold.*
2. *$\mathsf{max}(\mathfrak{P}) \models \phi$ implies $\mathfrak{P} \models \phi$. The reverse implication does not hold.*
3. *$\mathsf{min}(\mathfrak{P}) \models \phi$ if and only if $\mathfrak{P} \models \phi$.*

Proof: The implication in (1) is immediate by induction on ϕ. A counterexample for the reverse implication is given by the formula $\phi = !\{a\}$ and the cathoristic model $\mathfrak{M} = (\{s,t\}, s \xrightarrow{a} t, \lambda), s)$ where $\lambda(s) = \{a,b,c\}$: clearly $\mathsf{forget}(\mathfrak{M}) \models \phi$, but $\mathfrak{M} \not\models \phi$.

The implication in (2) is immediate by induction on ϕ. To construct a counterexample for the reverse implication, assume that Σ is a strict superset of $\{a\}$ a. The formula $\phi = !\{a\}$ and the pure cathoristic model $\mathfrak{P} = (\{s,t\}, s \xrightarrow{a} t), s)$ satisfy $\mathfrak{P} \models \phi$, but clearly $\mathsf{max}(\mathfrak{P}) \not\models \phi$.

Finally, (3) is also straightforward by induction on ϕ.

□

A.3 Non-determinism and cathoristic models

Both, cathoristic models and pure cathoristic models must be deterministic. That is important for the incompatibility semantics. However, formally, the definition of satisfaction makes sense for non-deterministic models as well, pure or otherwise. Such models are important in the theory of concurrent processes. Many of the theorems of the precious section either hold directly, or with small modifications for non-deterministic models. The rules of inference in Figure 12 are sound except for [Determinism] which cannot hold in properly non-deterministic models. With this omission, they are also complete. Elementary equivalence on non-deterministic cathoristic models also coincides with mutual simulation, while elementary equivalence on non-deterministic pure cathoristic models is bisimilarity. The proofs of both facts follow those of Theorems 1 and 11, respectively. Compactness by translation can be shown following the proof in Section 8, except that the constraint ϕ_{det} is unnecessary.

We have experimented with a version of cathoristic logic in which the models are *non-deterministic* labelled-transition systems. Although non-determinism makes some of the constructions simpler, non-deterministic cathoristic logic is unable to express incompatibility properly. Consider, for example, the claim that Jack is married[19] to Jill In standard deterministic cathoristic logic this would be rendered as:

$$\langle jack \rangle \langle married \rangle (\langle jill \rangle \wedge !\{jill\})$$

There are three levels at which this claim can be denied. First, we can claim that Jack is married to someone else - Joan, say:

$$\langle jack \rangle \langle married \rangle (\langle joan \rangle \wedge !\{joan\})$$

Second, we can claim that Jack is unmarried (specifically, that being unmarried is Jack's only property):

$$\langle jack \rangle !\{unmarried\}$$

Third, we can claim that Jack does not exist at all. Bob and Jill, for example, are the only people in our domain:

$$!\{bob, jill\}$$

Now we can assert the same sentences in non-deterministic cathoristic logic, but they are *no longer incompatible with our original sentence*. In non-deterministic cathoristic logic, the following sentences are compatible (as long as there are two separate transitions labelled with *married*, or two separate transitions labelled with *jack*):

$$\langle jack \rangle \langle married \rangle (\langle jill \rangle \wedge !\{jill\}) \qquad \langle jack \rangle \langle married \rangle (\langle joan \rangle \wedge !\{joan\})$$

[19] We assume, in this discussion, that *married* is a many-to-one predicate. We assume that polygamy is one person *attempting* to marry two people (but failing to marry the second).

Similarly, the following sentences are fully compatible as long as there are two separate transitions labelled with $jack$:

$$\langle jack\rangle\langle married\rangle \qquad \langle jack\rangle!\{unmarried\}$$

Relatedly, non-deterministic cathoristic logic does not satisfy Brandom's incompatibility semantics property:

$$\phi \models \psi \text{ iff } \mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$$

To take a simple counter-example, $\langle a\rangle\langle b\rangle$ implies $\langle a\rangle$, but not conversely. But in non-deterministic cathoristic logic, the set of sentences incompatible with $\langle a\rangle\langle b\rangle$ is identical with the set of sentences incompatible with $\langle a\rangle$.

A.4 Semantic characterisation of elementary equivalence

In Section 6.1 we presented a semantic analysis of elementary equivalence, culminating in Theorem 1 which showed that elementary equivalence coincides with $\simeq$, the relation of mutual simulation of models. We shall now carry out a similar analysis for pure cathoristic models, and show that elementary equivalence coincides with bisimilarity, an important concept in process theory and modal logics [25]. Bisimilarity is strictly finer on non-deterministic transition systems than $\simeq$, and more sensitive to branching structure. In the rest of this section, we allow non-deterministic pure models, because the characterisation is more interesting that in the deterministic case.

Definition 34 A pure cathoristic model $(\mathcal{L}, s)$ is finitely branching if its underlying transition system $\mathcal{L}$ is finitely branching.

Definition 35 A binary relation $\mathcal{R}$ is a *bisimulation* between pure cathoristic models $\mathfrak{P}_i = (\mathcal{L}_i), s_i)$ for $i = 1, 2$ provided (1) $\mathcal{R}$ is a bisimulation between $\mathcal{L}_1$ and $\mathcal{L}_2$, and (2) $(s_1, s_2) \in \mathcal{R}$. We say $\mathfrak{P}_1$ and $\mathfrak{P}_2$ are *bisimilar*, written $\mathfrak{P}_1 \sim \mathfrak{P}_2$ if there is a bisimulation between $\mathfrak{P}_1$ and $\mathfrak{P}_2$.

Definition 36 The *theory* of $\mathfrak{P}$, written $\mathsf{Th}(\mathfrak{P})$, is the set $\{\phi \mid \mathfrak{P} \models \phi\}$.

Theorem 11 *Let $\mathfrak{P}$ and $\mathfrak{P}'$ be two finitely branching pure cathoristic models. Then: $\mathfrak{P} \sim \mathfrak{P}'$ if and only if $\mathsf{Th}(\mathfrak{P}) = \mathsf{Th}(\mathfrak{P}')$.*

Proof: Let $\mathfrak{P} = (\mathcal{L}, w)$ and $\mathfrak{P}' = (\mathcal{L}', w')$ be finitely branching, where $\mathcal{L} = (W, \rightarrow)$ and $(W', \rightarrow')$. We first show the left to right direction, so assume that $\mathfrak{P} \sim \mathfrak{P}'$.

The proof is by induction on formulae. The only case which differs from the standard Hennessy-Milner theorem is the case for $!A$, so this is the only case we shall consider. Assume $w \sim w'$ and $w \models !A$. We need to show $w' \models !A$. From the semantic clause for $!$, $w \models !A$ implies $\lambda(w) \subseteq A$. If $w \sim w'$, then $\lambda(w) = \lambda'(w')$. Therefore $\lambda'(w') \subseteq A$, and hence $w' \models !A$.

The proof for the other direction is more involved. For states $x \in W$ and $x' \in W$, we write

$$x \equiv x' \qquad \text{iff} \qquad \mathsf{Th}((\mathcal{L}, x)) = \mathsf{Th}((\mathcal{L}', x')).$$

We define the bisimilarity relation:

$$Z = \{(x, x') \in \mathcal{W} \times \mathcal{W}' \mid x \equiv x'\}$$

To prove $w \sim w'$, we need to show:

- $(w, w') \in Z$. This is immediate from the definition of Z.
- The relation Z respects the transition-restrictions: if $(x, x') \in Z$ then $\lambda(x) = \lambda'(x')$
- The forth condition: if $(x, x') \in Z$ and $x \xrightarrow{a} y$, then there exists a y' such that $x' \xrightarrow{a} y'$ and $(y, y') \in Z$.

– The back condition: if $(x, x') \in Z$ and $x' \xrightarrow{a} y'$, then there exists a y such that $x \xrightarrow{a} y$ and $(y, y') \in Z$.

To show that $(x, x') \in Z$ implies $\lambda(x) = \lambda'(x')$, we will argue by contraposition. Assume $\lambda(x) \neq \lambda'(x')$. Then either $\lambda'(x') \not\subseteq \lambda(x)$ or $\lambda(x) \not\subseteq \lambda'(x')$. If $\lambda'(x') \not\subseteq \lambda(x)$, then $x' \not\models !\lambda(x)$. But $x \models !\lambda(x)$, so x and x' satisfy different sets of propositions and are not equivalent. Similarly, if $\lambda(x) \not\subseteq \lambda'(x')$ then $x \not\models !\lambda'(x')$. But $x' \models !\lambda'(x')$, so again x and x' satisfy different sets of propositions and are not equivalent.

We will show the forth condition in detail. The back condition is very similar. To show the forth condition, assume that $x \xrightarrow{a} y$ and that $(x, x') \in Z$ (i.e. $x \equiv x'$). We need to show that $\exists y'$ such that $x' \xrightarrow{a} y'$ and $(y, y') \in Z$ (i.e. $y \equiv y'$).

Consider the set of y'_i such that $x' \xrightarrow{a} y'_i$. Since $x \xrightarrow{a} y$, $x \models \langle a \rangle \top$, and as $x \equiv x'$, $x' \models \langle a \rangle \top$, so we know this set is non-empty. Further, since $(\mathcal{W}', \rightarrow')$ is finitely-branching, there is only a finite set of such y'_i, so we can list them $y'_1, ..., y'_n$, where $n >= 1$.

Now, in the Hennessy-Milner theorem for Hennessy-Milner logic, the proof proceeds as follows: assume, for reductio, that of the $y'_1, ..., y'_n$, there is no y'_i such that $y \equiv y'_i$. Then, by the definition of $\equiv$, there must be formulae $\phi_1, ..., \phi_n$ such that for all i in 1 to n:

$$y'_i \models \phi_i \text{ and } y \not\models \phi_i$$

Now consider the formula:

$$[a](\phi_1 \vee ... \vee \phi_n)$$

As each $y'_i \models \phi_i$, $x' \models [a](\phi_1 \vee ... \vee \phi_n)$, but x does not satisfy this formula, as each ϕ_i is not satisfied at y. Since there is a formula which x and x' do not agree on, x and x' are not equivalent, contradicting our initial assumption.

But this proof cannot be used in cathoristic logic because it relies on a formula $[a](\phi_1 \vee ... \vee \phi_n)$ which cannot be expressed in cathoristic logic: Cathoristic logic does not include the box operator or disjunction, so this formula is ruled out on two accounts. But we can massage it into a form which is more amenable to cathoristic logic's expressive resources:

$$\begin{aligned} [a](\phi_1 \vee ... \vee \phi_n) &= \neg\langle a \rangle \neg(\phi_1 \vee ... \vee \phi_n) \\ &= \neg\langle a \rangle(\neg\phi_1 \wedge ... \wedge \neg\phi_n) \end{aligned}$$

Further, if the original formula $[a](\phi_1 \vee ... \vee \phi_n)$ is true in x' but not in x, then its negation will be true in x but not in x'. So we have the following formula, true in x but not in x':

$$\langle a \rangle(\neg\phi_1 \wedge ... \wedge \neg\phi_n)$$

The reason for massaging the formula in this way is so we can express it in cathoristic logic (which does not have the box operator or disjunction). At this moment, the revised formula is *still* outside cathoristic logic because it uses negation. But we are almost there: the remaining negation is in innermost scope, and innermost scope negation can be simulated in cathoristic logic by the ! operator.

We are assuming, for reductio, that of the $y'_1, ..., y'_n$, there is no y'_i such that $y \equiv y'_i$. But in cathoristic logic without negation, we cannot assume that each y'_i has a formula ϕ_i which is satisfied by y'_i but not by y - it might instead be the other way round: ϕ_i may be satisfied by y but not by y'_i. So, without loss of generality, assume that $y'_1, ..., y'_m$ fail to satisfy formulae $\phi_1, ..., \phi_m$ which y does satisfy, and that $y'_{m+1}, ..., y'_n$ satisfy formulae $\phi_{m+1}, ..., \phi_n$ which y does not:

$$\begin{aligned} &y \models \phi_i \text{ and } y'_i \not\models \phi_i \quad i = 1 \text{ to } m \\ &y \not\models \phi_j \text{ and } y'_j \models \phi_j \quad j = m+1 \text{ to } n \end{aligned}$$

The formula we will use to distinguish between x and x' is:

$$\langle a \rangle(\bigwedge_{i=1}^{m} \phi_i \ \wedge \bigwedge_{j=m+1}^{n} \mathsf{neg}(y, \phi_j))$$

Here, neg is a meta-language function that, given a state y and a formula ϕ_j, returns a formula that is true in y but incompatible with ϕ_j. We will show that, since $y \not\models \phi_j$, it is always possible to construct $\mathsf{neg}(y, \phi_j)$ using the ! operator.

Consider the possible forms of ϕ_j:

- $\top$: this case cannot occur since all models satisfy $\top$.
- $\phi_1 \wedge \phi_2$: we know $y'_j \models \phi_1 \wedge \phi_2$ and $y \not\models \phi_1 \wedge \phi_2$. There are three possibilities:
  1. $y \not\models \phi_1$ and $y \models \phi_2$. In this case, $\mathsf{neg}(y, \phi_1 \wedge \phi_2) = \mathsf{neg}(y, \phi_1) \wedge \phi_2$.
  2. $y \models \phi_1$ and $y \not\models \phi_2$. In this case, $\mathsf{neg}(y, \phi_1 \wedge \phi_2) = \phi_1 \wedge \mathsf{neg}(y, \phi_2)$.
  3. $y \not\models \phi_1$ and $y \not\models \phi_2$. In this case, $\mathsf{neg}(y, \phi_1 \wedge \phi_2) = \mathsf{neg}(y, \phi_1) \wedge \mathsf{neg}(y, \phi_2)$.
- $!A$: if $y \not\models !A$ and $y'_j \models !A$, then there is an action $a \in \Sigma - A$ such that $y \xrightarrow{a} z$ for some z but there is no such z such that $y'_j \xrightarrow{a} z$. In this case, let $\mathsf{neg}(y, \phi_j) = \langle a \rangle \top$.
- $\langle a \rangle \phi$. There are two possibilities:
  1. $y \models \langle a \rangle \top$. In this case, $\mathsf{neg}(y, \langle a \rangle \phi) = \bigwedge_{y \xrightarrow{a} z} \langle a \rangle \mathsf{neg}(z, \phi)$.
  2. $y \not\models \langle a \rangle \top$. In this case, $\mathsf{neg}(y, \langle a \rangle \phi) = !\{b \mid \exists z. y \xrightarrow{b} z\}$. This set of bs is finite since we are assuming the transition system is finitely-branching.

□

Fig. 16: Worked example of neg. Note that the transition system on the left is non-deterministic.

We continue with a worked example of neg. Consider y and y'_j as in Figure 16. One formula that is true in y'_j but not in y is

$$\langle a \rangle (\langle b \rangle \top \wedge \langle c \rangle \top)$$

Now:

$$\begin{aligned}
&\mathsf{neg}(y, \langle a \rangle (\langle b \rangle \top \wedge \langle c \rangle \top)) \\
&\quad = \bigwedge_{y \xrightarrow{a} z} \langle a \rangle \mathsf{neg}(z, \langle b \rangle \top \wedge \langle c \rangle \top) \\
&\quad = \langle a \rangle \mathsf{neg}(z_1, \langle b \rangle \top \wedge \langle c \rangle \top) \wedge \langle a \rangle \mathsf{neg}(z_2, \langle b \rangle \top \wedge \langle c \rangle \top) \\
&\quad = \langle a \rangle (\langle b \rangle \top \wedge \mathsf{neg}(z_1, \langle c \rangle \top)) \wedge \langle a \rangle \mathsf{neg}(z_2, \langle b \rangle \top \wedge \langle c \rangle \top) \\
&\quad = \langle a \rangle (\langle b \rangle \top \wedge \mathsf{neg}(z_1, \langle c \rangle \top)) \wedge \langle a \rangle (\mathsf{neg}(z_2, \langle b \rangle \top) \wedge \langle c \rangle \top) \\
&\quad = \langle a \rangle (\langle b \rangle \top \wedge !\{b\}) \wedge \langle a \rangle (\mathsf{neg}(z_2, \langle b \rangle \top) \wedge \langle c \rangle \top) \\
&\quad = \langle a \rangle (\langle b \rangle \top \wedge !\{b\}) \wedge \langle a \rangle (!\{c\} \wedge \langle c \rangle \top)
\end{aligned}$$

The resulting formula is true in y but not in y'_j.

B Omitted proofs

B.1 Proof of Lemma 5

If $\mathfrak{M} \models \phi$ then $\mathfrak{M} \preceq \mathsf{simpl}(\phi)$.

Proof: We shall show $\mathsf{Th}(\mathsf{simpl}(\phi)) \subseteq \mathsf{Th}(\mathfrak{M})$. The desired result will then follow by applying Theorem 1. We shall show that

$$\text{If } \mathfrak{M} \models \phi \text{ then } \mathsf{Th}(\mathsf{simpl}(\phi)) \subseteq \mathsf{Th}(\mathfrak{M})$$

by induction on ϕ. In all the cases below, let $\mathsf{simpl}(\phi) = (\mathcal{L}, w)$ and let $\mathfrak{M} = (\mathcal{L}', w')$. The case where $\phi = \top$ is trivial. Next, assume $\phi = \langle a \rangle \psi$. We know $\mathfrak{M} \models \langle a \rangle \psi$ and need to show that $\mathsf{Th}(\mathsf{simpl}(\langle a \rangle \psi)) \subseteq \mathsf{Th}(\mathfrak{M})$. Since $(\mathcal{L}', w') \models \langle a \rangle \psi$, there is an x' such that $w' \xrightarrow{a} x'$ and $(\mathcal{L}', x') \models \psi$. Now from the definition of $\mathsf{simpl}()$, $\mathsf{simpl}(\langle a \rangle \psi)$ is a model combining $\mathsf{simpl}(\psi)$ with a new state w not appearing in $\mathsf{simpl}(\psi)$ with an arrow $w \xrightarrow{a} x$ (where x is the start state in $\mathsf{simpl}(\psi)$), and $\lambda(w) = \Sigma$. Consider any sentence ξ such that $\mathsf{simpl}(\langle a \rangle \psi) \models \xi$. Given the construction of $\mathsf{simpl}(\langle a \rangle \psi)$, ξ must be a conjunction of $\top$ and formulae of the form $\langle a \rangle \tau$. In the first case, $(\mathcal{L}', x')$ satisfies $\top$; in the second case, $(\mathcal{L}', x') \models \tau$ by the induction hypothesis and hence $(\mathcal{L}', w') \models \langle a \rangle \tau$.

Next, consider the case where $\phi = !A$, for some finite set $A \subset \Sigma$. From the definition of $\mathsf{simpl}()$, $\mathsf{simpl}(!A)$ is a model with one state s, no transitions, with $\lambda(s) = A$. Now the only formulae that are true in $\mathsf{simpl}(!A)$ are conjunctions of $\top$ and $!B$, for supersets $B \supseteq A$. If $\mathfrak{M} \models !A$ then by the semantic clause for $!$, $\lambda'(w') \subseteq A$, hence $\mathfrak{M}$ models all the formulae that are true in $\mathsf{simpl}(!A)$.

Finally, consider the case where $\phi = \psi_1 \wedge \psi_2$. Assume $\mathfrak{M} \models \psi_1$ and $\mathfrak{M} \models \psi_2$. We assume, by the induction hypothesis that $\mathsf{Th}(\mathsf{simpl}(\psi_1)) \subseteq \mathsf{Th}(\mathfrak{M})$ and $\mathsf{Th}(\mathsf{simpl}(\psi_2)) \subseteq \mathsf{Th}(\mathfrak{M})$. We need to show that $\mathsf{Th}(\mathsf{simpl}(\psi_1 \wedge \psi_2)) \subseteq \mathsf{Th}(\mathfrak{M})$. By the definition of $\mathsf{simpl}()$, $\mathsf{simpl}(\psi_1 \wedge \psi_2) = \mathsf{simpl}(\psi_1) \sqcap \mathsf{simpl}(\psi_2)$. If $\mathsf{simpl}(\psi_1)$ and $\mathsf{simpl}(\psi_2)$ are $\mathsf{inconsistent}$ (see the definition of $\mathsf{inconsistent}$ in Section 6.4) then $\mathfrak{M} = \bot$. In this case, $\mathsf{Th}(\mathsf{simpl}(\psi_1) \wedge \mathsf{simpl}(\psi_2)) \subseteq \mathsf{Th}(\bot)$. If, on the other hand, $\mathsf{simpl}(\psi_1)$ and $\mathsf{simpl}(\psi_2)$ are not $\mathsf{inconsistent}$, we shall show that $\mathsf{Th}(\mathsf{simpl}(\psi_1 \wedge \psi_2)) \subseteq \mathsf{Th}(\mathfrak{M})$ by reductio. Assume a formula ξ such that $\mathsf{simpl}(\psi_1 \wedge \psi_2) \models \xi$ but $\mathfrak{M} \not\models \xi$. Now $\xi \neq \top$ because all models satisfy $\top$. ξ cannot be of the form $\langle a \rangle \tau$ because, by the construction of merge (see Section 6.4), all transitions in $\mathsf{simpl}(\psi_1 \wedge \psi_2)$ are transitions from $\mathsf{simpl}(\psi_1)$ or $\mathsf{simpl}(\psi_2)$ and we know from the inductive hypothesis that $\mathsf{Th}(\mathsf{simpl}(\psi_1)) \subseteq \mathsf{Th}(\mathfrak{M})$ and $\mathsf{Th}(\mathsf{simpl}(\psi_2)) \subseteq \mathsf{Th}(\mathfrak{M})$. ξ cannot be $!A$ for some $A \subset \Sigma$, because, from the construction of merge, all state-labellings in $\mathsf{simpl}(\psi_1 \wedge \psi_2)$ are no more specific than the corresponding state-labellings in $\mathsf{simpl}(\psi_1)$ and $\mathsf{simpl}(\psi_2)$, and we know from the inductive hypothesis that $\mathsf{Th}(\mathsf{simpl}(\psi_1)) \subseteq \mathsf{Th}(\mathfrak{M})$ and $\mathsf{Th}(\mathsf{simpl}(\psi_2)) \subseteq \mathsf{Th}(\mathfrak{M})$. Finally, ξ cannot be $\xi_1 \wedge xi_2$ because the same argument applies to xi_1 and xi_2 individually. We have exhausted the possible forms of ξ, so conclude that there is no formula ξ such that $\mathsf{simpl}(\psi_1 \wedge \psi_2) \models \xi$ but $\mathfrak{M} \not\models \xi$. Hence $\mathsf{Th}(\mathsf{simpl}(\psi_1 \wedge \psi_2)) \subseteq \mathsf{Th}(\mathfrak{M})$. □

B.2 Proof of Lemma 6

If $\phi \models \psi$ then $\mathsf{simpl}(\phi) \preceq \mathsf{simpl}(\psi)$

Proof: By Theorem 1, $\mathsf{simpl}(\phi) \preceq \mathsf{simpl}(\psi)$ iff $\mathsf{Th}(\mathsf{simpl}(\psi)) \subseteq \mathsf{Th}(\mathsf{simpl}(\phi))$. Assume $\phi \models \psi$, and assume $\xi \in \mathsf{Th}(\mathsf{simpl}(\psi))$. We must show $\xi \in \mathsf{Th}(\mathsf{simpl}(\phi))$. Now $\mathsf{simpl}()$ is constructed so that:

$$\mathsf{simpl}(\psi) = \bigsqcup \{\mathfrak{M} \mid \mathfrak{M} \models \psi\}$$

So $\xi \in \mathsf{Th}(\mathsf{simpl}(\psi))$ iff for all models $\mathfrak{M}$, $\mathfrak{M} \models \psi$ implies $\mathfrak{M} \models \xi$. We must show that $\mathfrak{M} \models \phi$ implies $\mathfrak{M} \models \xi$ for all models $\mathfrak{M}$. Assume $\mathfrak{M} \models \phi$. Then since $\phi \models \psi$, $\mathfrak{M} \models \psi$. But since $\xi \in \mathsf{Th}(\mathsf{simpl}(\psi))$, $\mathfrak{M} \models \xi$ also.

□

B.3 Proof of Lemma 7

If $\mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$ then $\mathcal{J}(\mathsf{simpl}(\psi)) \subseteq \mathcal{J}(\mathsf{simpl}(\phi))$

Proof: Assume $\mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$ and $\mathfrak{M} \sqcap \mathsf{simpl}(\psi) = \bot$. We need to show $\mathfrak{M} \sqcap \mathsf{simpl}(\phi) = \bot$. If $\mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$ then for all formulae ξ, if $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\psi) = \bot$ then $\mathsf{simpl}(\xi) \sqcap \mathsf{simpl}(\phi) = \bot$. Let ξ be $\mathsf{char}(\mathfrak{M})$. Given that $\mathfrak{M} \sqcap \mathsf{simpl}(\psi) = \bot$ and $\mathsf{simpl}(\mathsf{char}(\mathfrak{M})) \preceq \mathfrak{M}$, $\mathsf{simpl}(\mathsf{char}(\mathfrak{M})) \sqcap \mathsf{simpl}(\psi) = \bot$. Then as $\mathcal{I}(\psi) \subseteq \mathcal{I}(\phi)$, $\mathsf{simpl}(\mathsf{char}(\mathfrak{M})) \sqcap \mathsf{simpl}(\phi) = \bot$. Now as $\mathfrak{M} \preceq \mathsf{simpl}(\mathsf{char}(\mathfrak{M}))$, $\mathfrak{M} \sqcap \mathsf{simpl}(\phi) = \bot$.

□